SHIELDWORKZ

FulcrumSec

Comprehensive Threat Actor Intelligence Dossier

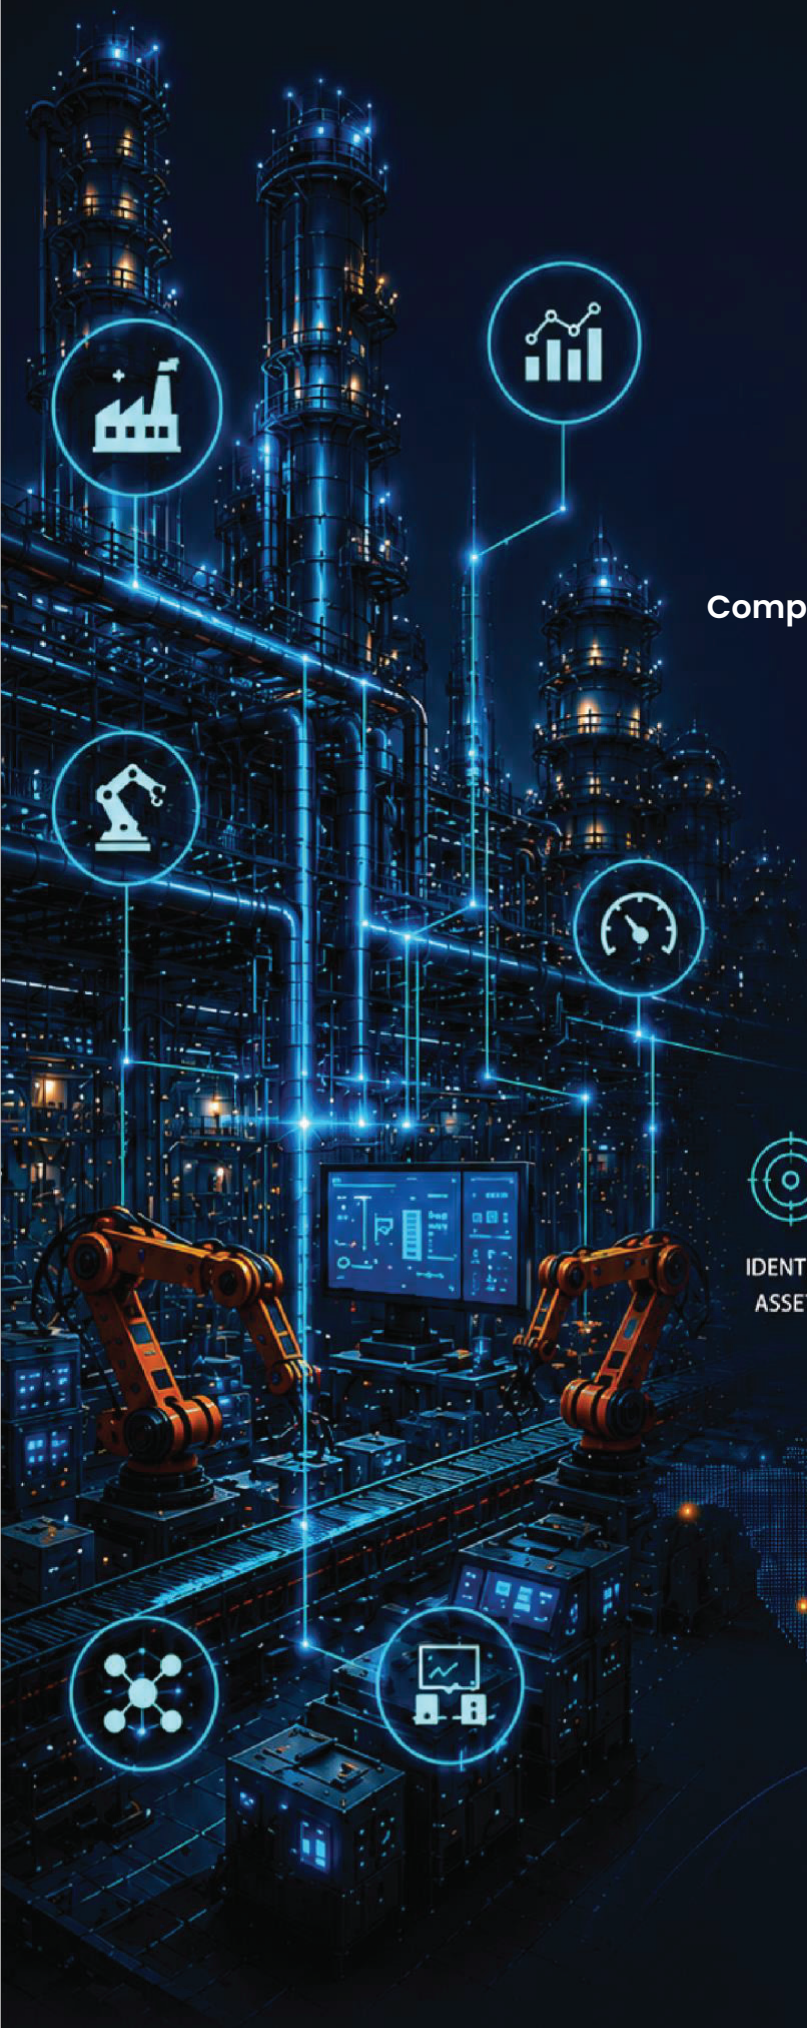

IDENTIFY ASSETS | ASSESS RISKS | EVALUATE CONTROLS | TREAT & PRIORITIZE RISKS | MONITOR & IMPROVE

This page has been intentionally left blank

Table of Contents

Table of Contents 3

How to Read This Dossier 6

1. Executive Threat Assessment 7

Why CISOs Should Care 8

CISO Risk Assessment 8

2. FulcrumSec Identity and Attribution 9

Name, Aliases, and Online Identities 9

Emergence and Evolution 9

Relationship to Ransomware and Extortion Ecosystems 10

3. Timeline 11

4. Manchester Airports Group Incident — Case Study 14

4.1 What MAG Officially Disclosed 14

4.2 What FulcrumSec Claimed 14

4.3 Technical Detail: Initial Access, Progression, Persistence, Lateral Movement ... 15

4.4 Operational and Customer Impact 16

4.5 Recovery and Containment 16

4.6 Attribution Confidence Assessment 16

4.7 Claim / Evidence / Source / Confidence Table 17

4.8 Lessons for Airports and Critical-Infrastructure Operators 18

5. Targeting Profile 19

Sector Analysis 19

Confirmed / Suspected / Claimed Targets and Strategic Pattern 19

6. Tactics, Techniques and Procedures (TTPs) — MITRE ATT&CK Mapping ... 21

7. Attack Playbook — Lifecycle 23

Reconnaissance 23
Initial Access 23
Execution 23
Persistence 23
Privilege Escalation 23
Credential Access 23
Discovery 24
Lateral Movement 24

Collection 24
Exfiltration 24
Impact 24
Extortion 24

8. Tools, Malware and Infrastructure 25

9. Initial Access Analysis 27
Door 1 — Exposed/Hardcoded Credentials 27
Door 2 — Exploitation of a Known, Patched Vulnerability 27
Door 3 — Misconfigured / Open Cloud Storage 27
Vectors With No Evidence 27

10. Identity, Credentials and Privilege Escalation 28
Defensive Controls and Telemetry 28

11. Lateral Movement and Internal Reconnaissance 29
Detection and Defensive Focus 29

12. Data Theft and Extortion 30
Extortion Model 30
Victim Communications and Pressure Tactics 30
Claimed Data Volumes and Types 30
Primary Focus Assessment 31

13. Infrastructure and Operational Security 32

14. Detection Opportunities — FulcrumSec Detection Matrix 33

15. IOC Catalogue 35

16. Defensive Strategy 36
Identity 36
Endpoint 36
Network 36
Email 36
Cloud 36
Data 37
Vulnerability Management 37
Incident Response 37

17. OT / ICS / Critical Infrastructure Implications 38
Observed OT Activity 38
Indirect OT Exposure 38
Plausible Attack Paths (Analytical, Not Observed) 38
Unsupported Speculation (Explicitly Rejected) 39
OT Defensive Considerations (Evidence-Based, Precautionary) 39

18. Sector-Specific Risk *40*

19. Threat Hunting Playbook *42*

Hypothesis 1 — Exposed secrets in public-facing code 42

Hypothesis 2 — Bulk repository or data-platform access from a single credential 42

Hypothesis 3 — Unpatched CVE-2025-55182 ("React2Shell") exposure 43

Hypothesis 4 — Sustained low-volume anomalous egress from a service identity 43

20. Incident Response Playbook *44*

Decision Points 45

21. Early Warning Indicators *46*

Actor-Specific Indicators (FulcrumSec) 46

Generic Ransomware/Extortion Indicators (Not FulcrumSec-Specific) 46

22. Unknowns and Intelligence Gaps *47*

What We Still Do Not Know 47

23. Myth vs Fact *49*

24. Intelligence Confidence Assessment *50*

25. CISO Executive Takeaways *51*

26. CISO Action Checklist *53*

Immediate — 0–7 Days 53

Near Term — 7–30 Days 53

Medium Term — 30–90 Days 53

Strategic — 90+ Days 53

27. Appendix *55*

27.1 Glossary 55

27.2 Acronyms 55

27.3 Confidence Methodology 56

27.4 Attribution Methodology 56

About Shieldworkz 57

How to Read This Dossier

This document distinguishes rigorously between four categories of information, marked throughout with colour-coded tags:

- **[FACT]** Directly established by reliable, corroborated evidence (official disclosures, independently validated technical detail, regulatory filings).
- **[ASSESSMENT]** A reasoned analytical conclusion drawn from available evidence by Shieldworkz CTI or the cited third-party researchers. Assessments are labelled with a confidence level (High / Moderate / Low).
- **[CLAIM]** An allegation made by FulcrumSec, a victim organisation, or another source that has not been independently established. Claims originating from the threat actor itself are additionally marked "Threat-actor claim — independently unverified."
- **[UNKNOWN]** Insufficient reliable public evidence exists. Where an entire sub-topic is unknown, this dossier states "No reliable public evidence identified" rather than speculating.

No aliases, victims, malware, infrastructure, techniques, or attribution have been invented for completeness. Where a required section of the source research brief cannot be populated with genuine evidence, that gap is stated explicitly rather than filled. All dates are given with their source and, where the underlying event date is uncertain, that uncertainty is flagged. This dossier reflects open-source and vendor threat-intelligence reporting available up to 2 September 2026; given FulcrumSec's ongoing activity and the Manchester Airports Group (MAG) incident's active, developing status, several assessments in this dossier are explicitly preliminary and may change as further information emerges.

1. Executive Threat Assessment

FulcrumSec is a financially motivated, data-theft-and-extortion group that has been publicly tracked by multiple independent threat-intelligence vendors since it first appeared in reporting around September–October 2025.

[FACT] FulcrumSec is a financially motivated cyber-extortion group first documented in reporting from approximately September–October 2025, also referred to in vendor reporting as "The Threat Thespians." It does not deploy ransomware/encryption; its model is data theft followed by extortion (a "steal-and-squeeze" model), corroborated independently by Sysdig, MOXFIVE, Dataminr, and Mallory.ai.

[ASSESSMENT] FulcrumSec is best characterised as a "cloud-native" or "identity-based" extortion actor rather than a classic ransomware operator: it consistently gains access through exposed credentials/API keys, cloud misconfigurations, and (in at least one case) an unpatched web application vulnerability, then exfiltrates data using legitimate cloud/SaaS tooling rather than custom malware (assessed High confidence, based on convergent reporting from Sysdig, MOXFIVE, and Mallory.ai across more than a dozen claimed intrusions).

[ASSESSMENT] Geographic focus is not narrowly regional: claimed/confirmed victims span the United States, the United Kingdom, Denmark, Australia, and Singapore, with reporting indicating most named victims are headquartered in the US (Sysdig, June 2026, citing roughly 25 claimed victims across 11 countries). The August 2026 MAG claim, if accurate, would mark a notable expansion into UK critical-transport infrastructure.

[ASSESSMENT] Sector focus is broad and opportunistic rather than narrowly strategic: technology, professional/engineering services, healthcare, pharmaceuticals, financial services/fintech, education, electronics manufacturing, and (per the FulcrumSec claim addressed in Section 4) transport/aviation. MOXFIVE noted the group had stated in early victim communications that it would not target healthcare organisations, yet healthcare victims (e.g., Lena Health) subsequently appeared on its leak site — an inconsistency between stated principles and observed conduct (MOXFIVE, May 2026).

[CLAIM] FulcrumSec has publicly listed, named, or claimed responsibility for intrusions against roughly 21–25+ organisations since late 2025, per its own leak-site postings and multiple vendor trackers (MOXFIVE, Sysdig, BreachNews, ransomware.live). The precise, current count is a moving target and vendors do not fully agree on totals, in part because some named organisations dispute or have not confirmed compromise.

[ASSESSMENT] Known capabilities are consistent with a technically competent but not highly novel operator: proficiency in cloud/SaaS environments (AWS, Azure, GCP, Databricks, MongoDB Atlas, GitHub, Azure DevOps, Hugging Face), credential/secrets-hunting discipline, patient/extended dwell time (reported as ranging from weeks to several months), narrative-control and PR-style extortion tactics (media kits sent to outlets such as vx-underground and BleepingComputer), and — per Mallory.ai reporting — the use of large language models to analyse stolen data and draft negotiation messages. No custom malware family or zero-day exploit has been publicly and reliably attributed to the group; its one documented CVE usage (CVE-2025-55182 / "React2Shell") is a known, patched vulnerability, not a zero-day (assessed Moderate–High confidence).

[UNKNOWN] The group's organisational structure, size, exact membership, nationality/location of operators, and any state or criminal-syndicate affiliation. No reliable public attribution to individuals or to a nation-state exists as of this writing.

Why CISOs Should Care

FulcrumSec exemplifies a broader and accelerating trend: encryption-free, identity- and cloud-based data extortion that bypasses most ransomware-era detection investment (EDR tuned for encryption behaviour, backup-focused resilience planning) and instead exploits basic credential hygiene and cloud configuration failures. Its claimed August 2026 intrusion into Manchester Airports Group — while limited, per MAG's own disclosure, to customer marketing/booking data rather than operational or safety systems — illustrates that critical-infrastructure operators are exposed through the same ordinary, internet-facing marketing and customer-engagement stack used by any consumer business, and that a single exposed third-party API credential can produce a breach notification affecting millions of people.

CISO Risk Assessment

Dimension	Assessment
Threat Level	Moderate-to-High for organisations with cloud-hosted customer data, exposed API/SaaS credentials, or hardcoded secrets in client-side code or repositories. Not currently assessed as a threat to OT/ICS or safety-critical systems (no evidence identified — see Section 17).
Confidence Level	Moderate. The group's general TTPs and victim history are corroborated by multiple independent CTI vendors (High confidence on modus operandi). The MAG attribution specifically is Moderate confidence only — see Section 4 Attribution table.
Primary Risks	Exposed/hardcoded API keys and tokens in client-side JavaScript and source repositories; unrotated cloud/SaaS credentials; misconfigured cloud storage; unpatched internet-facing web applications; reputational and regulatory fallout from large-scale customer-data exposure; secondary phishing/smishing campaigns built on exfiltrated contact and travel data.
Most Exposed Sectors	Technology, professional/engineering services, healthcare & pharmaceuticals, financial services/fintech, education, electronics/manufacturing distribution, and — per the claimed MAG intrusion — transport/aviation customer-facing platforms.
Recommended Immediate Actions	(1) Audit all client-side JavaScript/mobile app bundles for embedded API keys, especially marketing/CX platforms (Iterable, Marketo, Qualtrics and similar); (2) enforce mandatory credential rotation and least-privilege scoping for third-party SaaS integrations; (3) run secrets-scanning across Git history, CI/CD logs, and cloud storage; (4) validate that CVE-2025-55182 ("React2Shell") is patched on all internet-facing React deployments; (5) review cloud storage buckets and dev/staging subdomains for public exposure.

2. FulcrumSec Identity and Attribution

Name, Aliases, and Online Identities

[FACT] Primary name: "FulcrumSec." Reported alias: "The Threat Thespians" (used consistently across Sysdig, Dataminr, and Wasteland.me reporting).

[FACT] The group operates public-facing extortion/leak infrastructure including a clearnet site (reported domain fulcrumsec[.]net, hosting an "Index of /Shame" victim-shaming page) and a Tor (.onion) leak site referenced across multiple victim disclosures (e.g., the Arup Group case documented by researcher "BushidoToken," June 2026). A further clearnet-style domain (fulcrumsec[.]vg) has also been observed hosting campaign/victim listings and contact details, per open-source tracking.

[FACT] Contact channels reported in threat-intelligence write-ups include named e-mail addresses associated with the group's domain (e.g., a "threatspians@" mailbox referenced by Dataminr, March 2026), a Session messenger ID, and a Tox ID published on the group's site listing, per direct observation recorded in vendor reporting.

[CLAIM] FulcrumSec has, on at least one occasion, proactively contacted independent researchers/outlets (vx-underground, BleepingComputer) with "media kits" and data samples to substantiate and publicise its claims — behaviour multiple vendors (MOXFIVE, BleepingComputer) describe as active narrative management.

[UNKNOWN] No reliable public evidence identifies specific individual operators, nationality, or physical location of FulcrumSec members. No court filing, indictment, or law-enforcement attribution naming individuals tied to the "FulcrumSec" brand was identified in this research.

Emergence and Evolution

[FACT] Multiple independent vendors place FulcrumSec's first documented activity in September–October 2025, with its first confirmed/first-reported breach being the Avnet (electronics distributor) intrusion, reported around 7 October 2025 (Dataminr; Sysdig; TechNadu).

[ASSESSMENT] The group's public profile escalated significantly through H1 2026: from a single-victim posting in late 2025 (per Joe Shenouda's December 2025 write-up, which recorded "Confirmed Victims: One" — Avnet) to a dedicated leak site claiming 21 victims and two active campaigns by May 2026 (MOXFIVE), to roughly 25 claimed victims across 11 countries by June 2026 (Sysdig), and a headline-grabbing $25 million extortion attempt against Novo Nordisk in June 2026. This pattern indicates rapid scaling of operational tempo and public visibility over an approximately 9-month period (assessed High confidence based on convergent vendor timelines).

[CLAIM] A May 2026 MOXFIVE report notes FulcrumSec launched two internally branded "campaigns" — "The Hardcoded Horror Show" (targeting organisations with credentials exposed in code/CI-CD) and "Index of /Shame" (targeting organisations with publicly exposed storage/staging systems) — suggesting a degree of internal operational segmentation or marketing-style branding of intrusion types, per the group's own leak-site campaign pages.

Relationship to Ransomware and Extortion Ecosystems

[FACT] No ransomware/encryption payload has been publicly documented in any FulcrumSec intrusion. MOXFIVE states explicitly: "No ransomware binary or encryption behavior has been observed across known FulcrumSec operations."

[ASSESSMENT] Researchers (Sysdig, MOXFIVE) frame FulcrumSec within a broader 2025–2026 trend of "encryption-free" data extortion actors (comparable in operating model, though not confirmed to be affiliated with, groups such as LAPSUS$, Karakurt, and ShinyHunders/ShinyHunters-style identity-based crews) — a comparison offered as analytical context by researchers, not as evidence of shared membership or coordination (Moderate confidence; explicitly framed as a stylistic/operational analogy in the cited sources, not an attribution claim).

[CLAIM] During its June 2026 Novo Nordisk extortion, FulcrumSec publicly acknowledged that a second, separate actor calling itself "TheUSERS007" had also claimed a breach of Novo Nordisk systems around the same period (5–7 June 2026) using a self-described "venomware" tool and a separate $50 million demand. FulcrumSec stated it regarded the second group's claim as "potentially credible." Whether the two actors shared access, coordinated, or acted entirely independently is not established (Cybersecurity Insiders, gblock.app, June 2026).

[UNKNOWN] No reliable evidence of hacktivist ideology, political motivation, or state sponsorship was identified for FulcrumSec. All corroborated reporting characterises its motivation as financial.

3. Timeline

Dates below reflect either the date of the underlying event as disclosed by a victim/authority, or the date of first credible public reporting where the underlying event date is not independently fixed. Uncertain dates are marked accordingly.

Date	Event	Type	Source
~Sept 2025	FulcrumSec assessed to begin operations (earliest activity referenced in later vendor retrospectives).	ASSESSMENT	Sysdig; ransomware.live
7 Oct 2025	First confirmed/first-reported breach: Avnet (US electronics distributor), via misconfigured public-facing cloud storage linked to an internal sales tool, followed by use of a stolen OpenAI API key.	CLAIM (threat-actor account, partially corroborated by vx-underground contact)	Dataminr; Sysdig; vx-underground (X)
~10 Nov 2025	Claimed breach of US media company Blavity.	CLAIM	Dataminr; techtimes.com retrospective
19 Dec 2025	First detailed independent threat-intel write-up published (Joe Shenouda), recording "Confirmed Victims: One" (Avnet) and describing FulcrumSec as an emerging "data extortion-only" group.	FACT (publication date)	shenouda.nl
~26 Jan 2026	Claimed breach of US healthcare company Lena Health — notable because FulcrumSec had reportedly stated an intent not to target healthcare organisations.	CLAIM	Dataminr; MOXFIVE
Feb 2026	Claimed breach of LexisNexis via an unpatched public-facing web application vulnerability (later reporting identifies CVE-2025-55182, "React2Shell").	CLAIM, partially corroborated (LexisNexis confirmed a data breach per BleepingComputer reporting referenced by other vendors)	MOXFIVE; techtimes.com; bushidotoken.net
~Feb 2026	Claimed ~1.3 TB breach of Avnet including financial models and customer data (a larger, separate claim from the Oct 2025 initial-access report).	CLAIM	techtimes.com retrospective citing Sysdig
Feb 2026	Claimed ~300 GB breach of Australian fintech "youX" (formerly DRIVE IQ), reportedly affecting ~444,000 borrowers and ~800 broker organisations, following a breakdown in ransom negotiations; leak-site "Index of /Shame" section referenced an exposed MongoDB Atlas cluster.	CLAIM	Dataminr Cyber Intel Brief (5 Mar 2026); Sysdig
2026 (date uncertain)	Claimed breach of Global Schools Group/Foundation (Singapore-based education operator).	CLAIM	BreachNews; Mallory.ai; ransomware.live

Date	Event	Type	Source
10 May 2026	FulcrumSec publicly names UK engineering firm Arup Group on its Tor leak site, claiming ~700GB of GitHub repositories and ~2TB of Azure/AWS cloud and database data, including HS2- and Euston Station-related design files.	CLAIM (threat-actor leak-site posting; independent researcher documentation of the posting itself is FACT)	bushidotoken.net (10 Jun 2026 write-up); Mallory.ai
11 May 2026	MOXFIVE publishes detailed profile citing 21 claimed victims and two active named "campaigns" on FulcrumSec's leak site.	FACT (publication); underlying victim count is the actor's CLAIM	moxfive.com
11 June 2026	Novo Nordisk publicly discloses a cybersecurity incident, stating data was "copied externally without authorisation" from "a limited number of internal IT systems," including patient information.	FACT (victim disclosure)	pharmaphorum.com; SecurityWeek
13 June 2026 (approx., "two days later")	FulcrumSec claims responsibility for the Novo Nordisk breach via DataBreaches.net, stating access since March 2026 via a leaked GitHub personal access token, ~1.3TB exfiltrated, and a $25 million ransom demand; a second actor ("TheUSERS007") separately claims a related breach.	CLAIM (threat-actor)	DataBreaches.net (via pharmaphorum.com); SecurityWeek; Cybersecurity Insiders
Late June 2026	Novo Nordisk ransom demand refused; FulcrumSec begins exploring private sale of data and publishes successive data caches on its leak site.	CLAIM / partially corroborated (refusal reported by multiple outlets)	SecurityWeek; techtimes.com
26 June 2026	Sysdig publishes a consolidated "playbook" profile citing ~25 claimed victims across 11 countries and detailing FulcrumSec's three primary initial-access "doors" (exposed credentials, unpatched CVEs, misconfigured storage).	FACT (publication)	sysdig.com
24–27 Aug 2026	Manchester Airports Group (MAG) detects unauthorised access to customer-data systems (per MAG statement, detection preceded the 27 Aug public disclosure by approximately three days).	FACT (victim disclosure)	MAG public statement; techtimes.com
27 Aug 2026	MAG publicly discloses the breach, stating ~8.7 million customers affected across Manchester, London Stansted and East Midlands airports; no attacker named at time of disclosure.	FACT (victim disclosure)	MAG statement; BleepingComputer; Infosecurity Magazine; Help Net Security
~28–29 Aug 2026	MAG confirms to media that a ransom was demanded and that MAG refused to pay; ICO confirms receipt of a breach report and that it is assessing it; NCSC notified.	FACT (victim/regulator statements)	BBC (via tech-insider.org, cdotimes.com); ICO statement
30 Aug 2026	FulcrumSec tells BleepingComputer it was behind the MAG breach, claiming ~86GB stolen via exposed Iterable API credentials in client-side JavaScript;	CLAIM (threat-actor), partially corroborated by	BleepingComputer (30 Aug 2026); Hackmanac;

Date	Event	Type	Source
	BleepingComputer independently validates one sample record.	independent journalist verification of one record	SecurityAffairs; TechNadu
31 Aug–2 Sept 2026	Continued secondary reporting and analysis of the MAG/FulcrumSec claim; MAG has not confirmed FulcrumSec's identity or the specific Iterable/JavaScript attack-path account. No leak-site publication of the full alleged 86GB MAG dataset had been independently confirmed as of this writing.	ASSESSMENT / ongoing, preliminary	techtimes.com; securityaffairs.com; multiple secondary outlets

Source: Compiled from the sources cited in each row. This timeline will require updating as the MAG incident continues to develop.

4. Manchester Airports Group Incident — Case Study

4.1 What MAG Officially Disclosed

[FACT] MAG (operator of Manchester Airport, London Stansted Airport, and East Midlands Airport, the UK's largest airport operator by MAG's own description) publicly disclosed on 27 August 2026 that "a quantity of customer data has been obtained" by "an unauthorized third party," relating to car park, lounge, and Fast Track bookings, and in-airport Wi-Fi sign-ups across all three airports.

[FACT] MAG stated the exposed data fields include customers' email addresses, telephone numbers, postal codes, and vehicle registration numbers. MAG stated explicitly that neither MAG nor the affected system stored customers' bank details, payment card data, account passwords, or passport numbers, and that these categories were not exposed.

[FACT] MAG stated it "immediately contained the risk," engaged specialist external cybersecurity advisers, notified "the relevant authorities," and stressed that passenger safety and aviation security were not compromised and that airport operations (including parking services) continued to operate normally, with no cancelled flights or terminal disruption reported.

[FACT] As a precaution, MAG temporarily suspended its online "Manage My Booking" service and redirected affected customers to a telephone line; MAG stated it contacted affected customers directly, including those with upcoming bookings.

[FACT] Approximately 8.7 million customers are reported as affected, per widely cited UK media figures attributed to MAG; MAG itself, in its own public statement, did not disclose a specific number, and one outlet (BleepingComputer) noted local media had cited a slightly different, unconfirmed figure of up to 8.9 million.

[FACT] MAG subsequently confirmed to media (reported via the BBC and other outlets) that the attacker(s) demanded a ransom, and that MAG refused to pay.

[FACT] The UK Information Commissioner's Office (ICO) confirmed it received a breach report from MAG and is assessing it; no fine or formal enforcement outcome had been announced as of this writing. The National Cyber Security Centre (NCSC) was also notified, consistent with reporting obligations for operators touching critical national infrastructure.

[UNKNOWN] MAG has not published a technical root-cause statement describing the initial-access vector, and — as of this writing — has not publicly named FulcrumSec or any other group as the perpetrator. When BleepingComputer put FulcrumSec's specific claims (the 86GB figure and the Iterable API-credential access path) to MAG directly, a MAG spokesperson declined to address them, instead repeating a general statement about protective measures and customer notification.

4.2 What FulcrumSec Claimed

[CLAIM] On 30 August 2026, FulcrumSec told BleepingComputer, in direct correspondence, that it was responsible for the MAG breach. FulcrumSec claimed to have exfiltrated approximately 86 GB of data — considerably more, and more granular, than the categories MAG's own disclosure described. Labelled per this dossier's standard: Threat-actor claim — independently unverified, except where noted below.

[CLAIM] FulcrumSec claimed initial access was obtained via airport-specific API credentials for Iterable (a third-party customer-engagement/marketing-messaging platform used by MAG for functions such as booking confirmations and Wi-Fi sign-up emails), which the group said were left exposed in client-side (browser-executed) JavaScript — i.e., visible to anyone inspecting the website's source code, requiring no exploitation, credential theft from a third party, or malware. MAG has not confirmed this account. Threat-actor claim — independently unverified.

[CLAIM] FulcrumSec provided BleepingComputer with data samples including a roughly 21.5 GB "Manchester customer export" containing consolidated customer profiles combining identifiers with historical booking activity and marketing classifications, and claimed the full alleged dataset included purchase/booking references, airport and product selections, prices and discounts, booking status, parking dates and times, historical spending totals, IP addresses, approximate geolocation, device information, and customer-engagement metadata, plus nearly 200,000 records tied to travel scheduled for the remainder of 2026. Threat-actor claim — independently unverified beyond the single record described below.

[FACT] BleepingComputer independently validated one sample record supplied by FulcrumSec by comparing it against a real traveller's known Manchester Airport purchase history; the record accurately reflected that traveller's prior Fast Track purchases, booking and scheduled-arrival times, terminal used, amounts paid, purchase references, total spending, and apparent trip purpose. BleepingComputer explicitly stated it could not independently establish the overall claimed 86GB dataset size or verify the ~200,000 upcoming-travel-records claim.

[FACT] No payment card or banking information appeared in the samples BleepingComputer reviewed, consistent with MAG's own statement that such data was not held in the affected system.

[CLAIM] FulcrumSec told BleepingComputer it planned to publish the full alleged dataset if its demands were not met. As of this writing, no independent, corroborated confirmation of full public leak-site publication of the complete alleged MAG dataset was identified in this research.

4.3 Technical Detail: Initial Access, Progression, Persistence, Lateral Movement

[CLAIM] Initial access: per FulcrumSec's account only, exposed Iterable API credentials embedded in client-side JavaScript. This is consistent with FulcrumSec's documented modus operandi in prior claimed intrusions (see Sections 2, 8, 9) but has not been confirmed by MAG or by independent forensic reporting.

[UNKNOWN] Attack progression beyond initial access, persistence mechanisms, privilege escalation, and lateral movement within MAG's environment: No reliable public evidence identified. Neither MAG nor FulcrumSec has published a detailed technical account of activity between the alleged credential exposure and the claimed exfiltration of ~86GB of data. Given the claimed access vector (a marketing-platform API key), it is plausible that the described access path would not have required lateral movement within MAG's core network at all — i.e., API-level access to a third-party SaaS platform (Iterable) could itself expose the marketing/customer datasets described, without traversing MAG's internal infrastructure. This is an inference about what the claimed access method would technically permit, not a confirmed account of what occurred, and is offered as ASSESSMENT (Low-Moderate confidence) rather than fact.

4.4 Operational and Customer Impact

[FACT] No disruption to flight operations, airport operations, or aviation safety systems has been reported by MAG, the NCSC, or any media outlet. MAG's parking services and general airport operations continued as normal; the only customer-facing operational change reported was the temporary suspension of the "Manage My Booking" online service.

[ASSESSMENT] The principal customer-facing risk is secondary/downstream: the combination of exposed email addresses, phone numbers, postcodes, and vehicle registrations (and, per FulcrumSec's unverified claim, more granular booking, spending, location, and device data) creates strong raw material for targeted phishing and smishing campaigns impersonating MAG or its airports, particularly against customers with genuine upcoming travel bookings, per commentary from multiple independent security vendors quoted in coverage (Check Point, Illumio).

4.5 Recovery and Containment

[FACT] MAG stated it "immediately contained the risk," engaged external specialist cybersecurity advisers, and restricted access to affected systems; MAG has not published further technical detail on remediation, credential rotation, or the current status of the Iterable integration.

4.6 Attribution Confidence Assessment

> **Attribution Rating: MODERATE CONFIDENCE (as of 2 September 2026)**
>
> FulcrumSec's claim of responsibility is plausible, internally consistent with the group's well-documented modus operandi (credential/API-key exposure as an initial-access vector; rapid bulk exfiltration from SaaS/cloud platforms; media-outreach-driven extortion), and partially corroborated by one independently validated data sample. However, MAG itself has not confirmed FulcrumSec's identity or its specific technical account, and no forensic, law-enforcement, or regulatory finding independently confirms the group's involvement. The rating reflects a real but unconfirmed and unfalsified claim, not an established fact.

Supporting and limiting evidence for this rating:

- Supporting: FulcrumSec proactively and specifically claimed the intrusion to a reputable outlet (BleepingComputer), providing verifiable sample data rather than a vague or copy-paste claim (a pattern consistent with the group's established "narrative control" behaviour documented by MOXFIVE in earlier cases).
- Supporting: The claimed access mechanism (a SaaS marketing-platform API credential exposed in client-side JavaScript) is technically consistent with, though not identical to, FulcrumSec's previously documented "Door One: hardcoded/exposed credentials" pattern (Sysdig, June 2026), and with its history of targeting exactly this class of customer-engagement/CRM-adjacent platforms.
- Supporting: One data sample was independently validated by BleepingComputer against real traveller records.
- Limiting: MAG has neither confirmed nor denied FulcrumSec's specific account when directly asked; MAG's public statements are attacker-agnostic.

- Limiting: The overall claimed data volume (86GB) and the ~200,000 upcoming-travel-record claim remain unverified by any independent party as of this writing.
- Limiting: No forensic incident-response report, law-enforcement statement, or regulator finding attributing the intrusion to FulcrumSec (or ruling it out) has been published.

[UNKNOWN] Whether the same actor is responsible for both the intrusion/initial access and any subsequent extortion/leak activity, or whether these could involve different parties (e.g., an initial-access broker separate from the extortion "brand"), is not established. FulcrumSec's pattern in other cases (Novo Nordisk, Arup) is to conduct intrusion, extortion negotiation, and leak-site publication itself, which would argue against a split-actor scenario here — but this is an inference by analogy (ASSESSMENT, Low-Moderate confidence), not direct evidence specific to the MAG case.

[UNKNOWN] No indication of impersonation (i.e., a third party falsely claiming credit under the FulcrumSec name) was identified, but this possibility cannot be excluded given the absence of independent forensic confirmation.

4.7 Claim / Evidence / Source / Confidence Table

Claim	Evidence	Source	Confidence	What Remains Unknown
A breach of MAG customer data occurred, affecting ~8.7M customers.	Direct victim disclosure; multiple independent media outlets.	MAG public statement, 27 Aug 2026	Confirmed (Fact)	Exact final affected-record count; whether the figure will be revised.
No payment/banking data was exposed.	Direct victim statement; consistent across all reviewed FulcrumSec-supplied samples.	MAG statement; BleepingComputer sample review	High	Whether any edge-case records (e.g., partial card metadata) exist outside reviewed samples.
FulcrumSec was responsible for the intrusion.	Threat-actor self-claim to a reputable outlet; one validated data sample.	FulcrumSec (via BleepingComputer, 30 Aug 2026)	Moderate	No MAG confirmation; no forensic/law-enforcement attribution.
~86GB of data was exfiltrated, including ~200,000 upcoming-travel records.	Threat-actor claim; not independently sized or verified.	FulcrumSec (via BleepingComputer)	Low-Moderate (Unverified)	True total volume; whether figures are inflated for extortion leverage.
Access was via exposed Iterable API credentials in client-side JavaScript.	Threat-actor technical account only; MAG declined to confirm or deny.	FulcrumSec (via BleepingComputer)	Low-Moderate (Unverified)	True initial-access vector; whether other vectors were also used.
MAG was asked to pay a ransom and refused.	Victim confirmation reported via BBC and corroborated across outlets.	MAG (via BBC); multiple secondary outlets	High	Amount demanded; identity of the party that made the demand (presumed but not confirmed to be FulcrumSec).

4.8 Lessons for Airports and Critical-Infrastructure Operators

1. Third-party marketing/CX SaaS integrations (e.g., customer-engagement platforms like Iterable) sit outside classical OT/critical-system security perimeters but can hold — or provide API-level access to — millions of customer records; they warrant the same credential-hygiene rigor as core infrastructure.
2. Segmentation between customer-facing digital services and safety/operational technology worked as intended in this case: MAG's own statements and independent reporting are consistent in finding no impact on flight operations or aviation safety systems. This is a positive control example worth preserving, not weakening, in incident response.
3. Attacker claims and victim disclosures can diverge substantially in scope; CISOs at peer organisations should not assume an official breach notice describes the full extent of exposure, and should plan customer communications and phishing-awareness messaging around the more conservative (larger) end of plausible exposure.
4. Rapid, transparent engagement with a researcher/outlet (as FulcrumSec did with BleepingComputer) is now a standard extortion-pressure tactic; incident-response and communications plans should anticipate media outreach by the attacker in parallel with, or ahead of, ransom negotiation.

5. Targeting Profile

Sector Analysis

Sector	Status
Aviation / Airports	Claimed (MAG, Aug 2026) — first and only aviation-sector claim identified to date.
Transport (broader)	Claimed (MAG only, as above).
Technology	Confirmed/Claimed — multiple named victims across vendor reporting (e.g., Avnet-adjacent tech operations; general sector inclusion per Sysdig, MOXFIVE).
Professional / Engineering Services	Claimed — Arup Group (May 2026), a major UK engineering consultancy.
Healthcare / Pharmaceuticals	Claimed — Novo Nordisk (June 2026), Lena Health (Jan 2026), Woundtech (per BreachNews listing, unconfirmed by independent breach coverage).
Financial Services / FinTech	Claimed — youX/DRIVE IQ (Feb 2026), LexisNexis Risk Solutions (Feb 2026), MyComplianceOffice.
Electronics Manufacturing / Distribution	Claimed/Confirmed — Avnet (Oct 2025 / expanded Feb 2026 claim).
Education	Claimed — Global Schools Group/Foundation (Singapore).
Media	Claimed — Blavity (Nov 2025).
Government	No reliable public evidence identified of a direct government-sector victim.
Energy	No reliable public evidence identified.
Other (misc. named, unconfirmed)	Stuf Storage, Hatica, ReFocus, Raptor Supplies, FashionZA, CrediElite, Rotary International, IMEVI, Interzero, SalesKido, ParkEngage, Nordstern Technologies/NCS — named on FulcrumSec leak-site/campaign pages per BreachNews tracking; BreachNews explicitly notes several of these "have not yet been the subject of public breach claims covered" independently, and their appearance "should not be interpreted as confirmation of compromise."

Confirmed / Suspected / Claimed Targets and Strategic Pattern

[FACT] Confirmed targets (i.e., victim organisation itself has publicly acknowledged a breach consistent with FulcrumSec's claim, even where the victim has not named FulcrumSec specifically): Novo Nordisk (confirmed unauthorised external data copying, 11 June 2026); Manchester Airports Group (confirmed customer-data breach, 27 Aug 2026, attacker unnamed by MAG).

[CLAIM] Claimed-but-not-independently-confirmed-by-victim targets: Avnet, Blavity, Lena Health, LexisNexis (partial third-party corroboration exists per BleepingComputer reporting

referenced by other vendors), youX, Arup Group, Global Schools Group, and the additional names listed on FulcrumSec's own leak site.

[ASSESSMENT] FulcrumSec's targeting pattern is best characterised as opportunistic-and-financially-motivated rather than strategically or ideologically selective: the common thread across victims is not sector or geography but the presence of an exploitable credential/configuration weakness (exposed API keys, unpatched CVE-2025-55182 deployments, or open cloud storage), consistent with the group's own leak-site framing of its campaigns around vulnerability classes ("Hardcoded Horror Show," "Index of /Shame") rather than around named sectors or countries (assessed High confidence, based on MOXFIVE and Sysdig's convergent characterisation).

[ASSESSMENT] Victim selection value appears influenced by data sensitivity and organisational revenue/profile (used to calculate ransom demands) rather than by a fixed sectoral doctrine — e.g., MOXFIVE and BreachNews both note listed targets "by revenue in the tens of billions of dollars," suggesting a preference for larger organisations once a technical foothold is found, but not a rule against smaller ones (Moderate confidence).

Strategic pattern summary: opportunistic and financially motivated, with technique-driven rather than sector-driven target selection; the inclusion of a UK critical-transport operator (MAG) in August 2026, if the claim is accurate, would represent an extension of this opportunistic model into critical infrastructure rather than evidence of a deliberate strategic pivot toward the transport sector specifically.

6. Tactics, Techniques and Procedures (TTPs) — MITRE ATT&CK Mapping

Only techniques with direct, cited evidentiary support (from FulcrumSec's own leak-site disclosures, or from independent vendor technical analysis of specific claimed intrusions) are listed. Techniques common to "ransomware groups generally" are not inferred onto FulcrumSec absent specific evidence.

ATT&CK ID / Technique	Evidence & Example	Confidence	Detection / Defensive Controls
T1190 — Exploit Public-Facing Application	Exploitation of CVE-2025-55182 ("React2Shell"), an unauthenticated RCE against unpatched React front-end applications, documented by MOXFIVE in at least one FulcrumSec intrusion (reported as the LexisNexis case) where the vulnerability reportedly went unpatched for 2+ months after public disclosure of active exploitation.	Moderate-High	Patch management SLAs for internet-facing frameworks; WAF rules for known RCE patterns; monitor for anomalous process execution on public web servers.
T1552.001 — Unsecured Credentials: Credentials in Files	Hardcoded GitHub personal access tokens, Azure container-registry credentials, and OpenAI/Iterable API keys found embedded in client-side JavaScript and source repositories across the Avnet, Novo Nordisk, Arup, and (claimed) MAG intrusions.	High	Secrets-scanning across Git history, CI/CD pipelines, and deployed front-end bundles; pre-commit hooks; mandatory credential rotation; DLP for repository exfiltration.
T1552.005 — Unsecured Credentials: Cloud Instance Metadata API	Cited by MOXFIVE as part of FulcrumSec's documented technique set (leak-site self-reporting and vendor analysis of cloud-focused intrusions).	Moderate	Restrict/require IMDSv2; monitor for anomalous metadata-service queries; cloud workload identity hardening.
T1580 — Cloud Infrastructure Discovery	Reconnaissance of victim cloud environments (AWS, Azure, GCP, Databricks) following initial credential compromise, cited by MOXFIVE and consistent across multiple claimed intrusions.	Moderate-High	Cloud-native detection (CloudTrail/Azure Activity Log/GCP Audit Log) for enumeration API calls from anomalous sources; CSPM alerting.
T1213 — Data from Information Repositories	Cloning of private GitHub/Azure DevOps repositories after initial token compromise (Novo Nordisk, Arup), yielding further embedded credentials and source code.	High	Repository access logging and anomaly detection; scoped, short-lived tokens rather than long-lived PATs; branch-protection and audit-log review.
T1530 — Data from Cloud Storage Object	Exfiltration from misconfigured/open cloud storage (S3, Azure Blob) and data platforms (MongoDB Atlas, Databricks) documented in the youX and other "Index of /Shame" cases.	High	Public-bucket/exposure scanning; least-privilege bucket policies; storage access logging with anomalous-download alerting.
T1567 — Exfiltration Over Web Service	Use of standard cloud/SaaS platform tooling for bulk data transfer rather than custom exfiltration malware, per MOXFIVE ("no	High	Egress monitoring/DLP tuned to cloud-storage and SaaS-API transfer volumes; anomalous

ATT&CK ID / Technique	Evidence & Example	Confidence	Detection / Defensive Controls
	ransomware binary or encryption behavior... Exfiltration conducted... using standard platform tools").		data-transfer-volume alerting per identity/service account.
T1585 / T1583 (media & narrative infrastructure — non-standard use)	Proactive outreach to independent researchers and media outlets (vx-underground, BleepingComputer) with pre-packaged "media kits" and data samples as an extortion-pressure tactic; not a conventional ATT&CK-mapped technique but a documented, recurring behavioural TTP.	High (behavioural, not ATT&CK-standard)	Incident-response and legal/communications playbooks should anticipate proactive attacker-to-media contact; monitor named breach-notification/journalist outlets and social media for early leakage of incident details.

[UNKNOWN] No reliable evidence of custom malware, C2 frameworks, credential-dumping tools (e.g., Mimikatz), living-off-the-land binaries beyond standard cloud/SaaS tooling, or classic Windows-domain lateral-movement techniques (T1021 variants, PsExec, WMI) was identified for FulcrumSec specifically. These are commonly seen in other ransomware/extortion groups but are not documented as part of FulcrumSec's observed toolkit, and this dossier does not infer their use.

7. Attack Playbook — Lifecycle

Mapped only where evidence exists. Stages without public evidence are marked accordingly rather than inferred.

Reconnaissance

[ASSESSMENT] Assessed to include automated scanning for exposed API keys/secrets in public repositories and client-side code, and for open cloud storage/staging systems ("Index of /Shame" campaign framing implies systematic scanning for this exposure class). Moderate confidence — inferred from campaign branding and outcome pattern rather than directly observed scanning telemetry.

Initial Access

[FACT] Three documented vectors, per MOXFIVE's June 2026 profile: (1) hardcoded/exposed credentials and API keys in client-side JavaScript, source repositories, or CI/CD logs; (2) exploitation of CVE-2025-55182 ("React2Shell") on unpatched public-facing React applications; (3) discovery of open/misconfigured cloud storage and staging environments requiring no exploitation. See Section 9 for detail.

Execution

[UNKNOWN] No reliable public evidence identified of a distinct "execution" stage involving payload delivery or code execution on victim endpoints beyond the initial web-application exploitation (T1190) itself. The group's model largely bypasses traditional endpoint execution by operating at the API/cloud-identity layer.

Persistence

[CLAIM] In the Novo Nordisk case, FulcrumSec claimed to have found additional credentials over an extended period (reportedly continuing even after the victim became aware of the breach), implying a persistence strategy based on credential re-discovery/rotation-evasion rather than classic host-based persistence mechanisms (e.g., scheduled tasks, registry run keys). Threat-actor claim — independently unverified in technical detail, though the extended timeline is broadly consistent with GitHub's and Azure's staggered detection timestamps reported by CybelAngel.

Privilege Escalation

[CLAIM] In the Novo Nordisk case, a JWT "alg:none" authentication-bypass vulnerability on a development-environment clinical API (COSMOS) was reportedly leveraged to escalate access, per CybelAngel's technical write-up drawing on FulcrumSec's own disclosed account. Threat-actor-sourced technical claim, independently described by a named vendor analysis rather than by the victim — Moderate confidence.

Credential Access

[FACT] Well documented across multiple cases: harvesting of additional API tokens, database credentials, and service-account passwords from cloned source-code repositories once an initial token provided repository access (Novo Nordisk, Arup).

Discovery

[ASSESSMENT] Cloud infrastructure discovery (enumeration of AWS/Azure/GCP resources, storage buckets, and data platforms) assessed as a standard stage based on the breadth of platforms reported compromised per victim (Novo Nordisk: GitHub, Azure DevOps, AWS, Hugging Face, Okta). Moderate-High confidence.

Lateral Movement

[UNKNOWN] No reliable public evidence identified of classic internal-network lateral movement (RDP, SMB, PsExec, WinRM, SSH pivoting). FulcrumSec's documented "lateral movement" is better described as credential-to-credential pivoting across cloud/SaaS platforms and repositories rather than movement across an internal Windows/AD network — see Section 11.

Collection

[FACT] Systematic collection of source code, database contents, employee/customer/patient records, financial models, proprietary research, and — per Mallory.ai — AI-assisted identification of "high-value relationships within stolen data" prior to extortion outreach.

Exfiltration

[FACT] Bulk exfiltration via standard cloud/SaaS transfer tooling (no custom exfiltration malware documented); volumes claimed have ranged from tens of gigabytes to 1.3+ terabytes across different victims. One vendor (tech-insider.org, citing FulcrumSec's own claim) states the group "deliberately throttled exfiltration to stay under behavioral thresholds" in at least one case — Threat-actor claim, independently unverified.

Impact

[FACT] No destructive or encryption impact documented in any case. Impact is limited to data confidentiality loss and its downstream reputational/regulatory/fraud consequences.

Extortion

[FACT] Consistent "steal-and-squeeze" extortion model: direct ransom negotiation with the victim; publication of data samples, internal communications, and (in at least one case) court-summons documents to pressure victims who downplay the breach; a dedicated public/dark-web leak site with victim-shaming pages; willingness to privately sell data rather than only leak it publicly when a sale is deemed more valuable; and proactive media outreach to amplify pressure. See Section 12 for full analysis.

8. Tools, Malware and Infrastructure

[FACT] No custom malware family or ransomware binary has been publicly and reliably attributed to FulcrumSec. MOXFIVE states explicitly this group has not been observed using ransomware/encryption tooling. All documented "tooling" consists of legitimate/standard cloud and developer platforms abused for access and exfiltration.

Name / Item	Purpose	Evidence	Confidence	Attribution
fulcrumsec[.]net (clearnet)	Public leak/shaming site, hosts "/shame" (Index of /Shame) victim-listing page.	Directly observed by multiple researchers across victim disclosures.	High	Directly attributed (group-operated domain)
fulcrumsec[.]vg (clearnet-style)	Victim/campaign listing page with contact emails; observed hosting the group's public campaign pages.	Directly observed in OSINT tracking.	Moderate-High	Directly attributed
Tor (.onion) leak site	Full stolen-data publication for victims who do not pay; referenced across multiple victim postings (e.g., "COMPLETE LEAKED AVNET DATA (.onion)").	Referenced in group's own site listings and third-party trackers.	High	Directly attributed
Session messenger ID / Tox ID	Encrypted negotiation/communication channel with victims and media.	Published on the group's own site listing.	High	Directly attributed
Named contact e-mail address(es) on fulcrumsec[.]net domain	Victim/media/researcher contact.	Cited by Dataminr (Mar 2026) and other trackers.	Moderate-High	Directly attributed
GitHub / Azure DevOps / AWS / GCP / Databricks / MongoDB Atlas / Hugging Face / Okta (legitimate platforms, abused)	Not group infrastructure — victim-hosted platforms exploited for access, discovery, and exfiltration via stolen credentials.	Consistently documented across Novo Nordisk, Arup, youX, and other claimed intrusions.	High	Associated (victim/legitimate infrastructure abused, not FulcrumSec-owned)
Stolen OpenAI API key (Avnet case)	Reported use of a victim's own stolen OpenAI credentials, "weaponized... against themselves" per Sysdig, reportedly to assist analysis of stolen data.	Sysdig (June 2026).	Moderate	Threat-actor use of victim-owned resource
LLM/AI tooling (general)	Reported use of large language models for analysing complex exfiltrated databases, identifying high-value data relationships, and	Mallory.ai analytical summary of multiple incidents.	Moderate	Assessed group practice, not a named product

Name / Item	Purpose	Evidence	Confidence	Attribution
	drafting negotiation messages.			

[UNKNOWN] No confirmed cryptocurrency wallet addresses, dedicated C2 IP infrastructure, bulletproof hosting provider, or phishing infrastructure specific to FulcrumSec was identified in this research. No reliable public evidence identified for these categories; readers should not assume their absence from public reporting means they do not exist — only that no corroborated public disclosure was located.

9. Initial Access Analysis

FulcrumSec's documented initial-access repertoire is narrow and consistent, per MOXFIVE's framing of "three doors":

Door 1 — Exposed/Hardcoded Credentials

[FACT] The dominant, best-evidenced vector. Examples: GitHub personal access tokens and Azure container-registry credentials embedded in client-side JavaScript on forgotten/dev subdomains (Novo Nordisk); a GitHub PAT hardcoded in JavaScript on a forgotten Arup subdomain; a stolen OpenAI API key at Avnet; airport-specific Iterable API credentials in client-side JavaScript (claimed, MAG). Confidence: High for the general pattern; Moderate for the MAG-specific instance pending independent confirmation.

Door 2 — Exploitation of a Known, Patched Vulnerability

[FACT] CVE-2025-55182 ("React2Shell"), an unauthenticated RCE affecting unpatched React front-end applications, documented in at least one case (reported as LexisNexis) where the victim had reportedly not patched for more than two months after active exploitation was publicly reported. This is the only specific CVE publicly and reliably associated with FulcrumSec activity — no zero-day use has been documented.

Door 3 — Misconfigured / Open Cloud Storage

[FACT] Publicly exposed cloud storage, staging servers, and databases requiring no credential theft or exploitation at all — the basis of the group's "Index of /Shame" campaign branding, documented in the youX case (unsecured MongoDB Atlas cluster) and referenced generally by MOXFIVE across "dozens of victims."

Vectors With No Evidence

- Phishing (email-based social engineering) — no reliable public evidence identified of FulcrumSec using phishing as an initial-access vector; the group's documented pattern is technical/credential-based, not human-targeted deception, though one earlier speculative write-up (Joe Shenouda, drawing an analogy to LAPSUS$/Karakurt) suggested SOCs "should hunt for" help-desk social engineering as a precaution — this is an analyst recommendation based on comparison to other groups, not evidence of FulcrumSec behaviour, and is explicitly marked here as such.
- VPN compromise, remote-access-tool abuse, supply-chain compromise, insider access, and zero-day exploitation — no reliable public evidence identified for any of these vectors in connection with FulcrumSec.
- Valid-accounts abuse via credential-stuffing or brute force — no reliable public evidence identified; the group's "valid credential" access has, in every documented case, derived from discovery of already-exposed secrets rather than from cracking or guessing.

10. Identity, Credentials and Privilege Escalation

[FACT] Credential theft (of already-exposed static secrets, not of live user sessions) is FulcrumSec's core and best-evidenced technique — see Sections 6 and 9.

[CLAIM] A single reported instance of authentication-bypass-based privilege escalation: an "alg:none" JWT vulnerability on a Novo Nordisk development-environment clinical API (COSMOS), allowing token acceptance without signature verification — per CybelAngel's analysis of FulcrumSec's disclosed account. This is a well-documented vulnerability class in the industry generally, and its application here is described by an independent vendor, giving it Moderate confidence, though it remains sourced from the threat actor's own account of the intrusion.

[UNKNOWN] No reliable public evidence identified of: token/session hijacking of live user sessions, MFA bypass techniques, Active Directory abuse (Kerberoasting, DCSync, Golden/Silver Ticket), or on-premises privileged-account compromise. FulcrumSec's documented activity operates almost entirely within cloud/SaaS/DevOps identity contexts rather than traditional Windows-domain identity infrastructure.

[ASSESSMENT] Cloud identity compromise (service-account and API-token abuse, rather than interactive human-identity compromise) is the group's primary "privilege escalation" pathway in practice: possession of one exposed token routinely yielded access to further, more privileged tokens and credentials stored in the same repositories or environments (documented in Novo Nordisk and Arup). Assessed High confidence as a general pattern.

Defensive Controls and Telemetry

- Enforce short-lived, scoped tokens for CI/CD, repository, and third-party SaaS integrations rather than long-lived personal access tokens; monitor and alert on PAT/API-key creation, use from new locations, and use outside expected service contexts.
- Deploy secrets-scanning (pre-commit, CI/CD pipeline, and periodic full-repository/history scans) with mandatory rotation on any detected exposure.
- Instrument cloud audit logs (CloudTrail, Azure Activity Log, GCP Audit Log, GitHub/Azure DevOps audit logs) for anomalous authentication and repository-cloning events, particularly bulk repository clones or unusual download volumes from a single token.
- For any development environment exposing authentication logic (e.g., JWT validation), ensure signature verification cannot be bypassed (reject `alg:none` and unsigned tokens) and treat "development"/"sandbox" environments as requiring production-grade authentication controls, since FulcrumSec has specifically targeted forgotten or assumed-private dev/sandbox subdomains.

11. Lateral Movement and Internal Reconnaissance

[UNKNOWN] No reliable public evidence identified of classic internal-network lateral movement techniques: RDP, SMB-based movement, PsExec, WinRM, SSH pivoting between hosts, or on-premises Active Directory reconnaissance/trust abuse. This dossier does not infer these techniques onto FulcrumSec despite their prevalence among other extortion actors, per the research brief's explicit instruction not to assume generic ransomware-group behaviour applies.

[FACT] FulcrumSec's documented "lateral movement" equivalent is cloud-to-cloud and platform-to-platform credential pivoting: a single exposed token in one system (e.g., a GitHub PAT) was used to access repositories, which yielded further credentials granting access to additional, distinct cloud platforms (Azure, AWS, Hugging Face, Okta in the Novo Nordisk case). This is a materially different risk model than classical internal-network lateral movement.

[CLAIM] MOXFIVE reported one instance in which "a single compromised host reached three organizations sharing one [cloud] account" — illustrating a cross-tenant/shared-account blast-radius risk distinct from within-organisation lateral movement. Sourced to MOXFIVE's analysis; the underlying incident details were not independently re-verified for this dossier.

Detection and Defensive Focus

- Monitor cloud IAM logs for a single identity/token accessing an unusual number of distinct services or platforms in a short window.
- Segment cloud accounts by business entity/environment so that compromise of one account or token cannot reach unrelated organisational data.
- Because FulcrumSec's movement is credential-pivot-based rather than network-based, traditional network-segmentation and NDR controls provide only partial protection; identity-centric detection (CIEM/cloud identity governance) is the primary relevant control.

12. Data Theft and Extortion

Extortion Model

[FACT] FulcrumSec operates a single, consistently observed extortion model that the group itself brands "steal and squeeze": data theft without encryption, followed by a direct ransom demand, followed — if unmet — by public leak-site posting, sample publication, and/or private sale of the stolen data. No ransomware/encryption component has ever been documented.

[FACT] This is a double-extortion-style model in effect (theft + publish/sell threat) but without the "first extortion" leg (encryption/operational disruption) that defines classic double extortion; it is more precisely a single-vector data-extortion model. No evidence of a "triple extortion" element (e.g., DDoS layered on top, or direct extortion of the victim's customers/patients) was identified as a standard practice, though the group has separately contacted or referenced third parties (e.g., outreach to researchers/media) as a pressure tactic rather than a distinct extortion channel against downstream victims.

Victim Communications and Pressure Tactics

- Direct ransom negotiation with the victim, with demands scaled to the target's apparent financial profile (e.g., $25 million demanded of Novo Nordisk; a demand was also made of MAG per BBC reporting, though the amount was not disclosed).
- Public "counter-messaging": if a victim publicly downplays a breach (e.g., claims data is "unreadable" or "unimportant"), FulcrumSec has responded by releasing plaintext samples to contradict the victim's framing (documented pattern per Dataminr's youX analysis).
- Publication of internal communications, court-summons documents, and executive names to increase reputational pressure (documented in the youX and other cases).
- Proactive outreach to independent researchers (vx-underground) and journalists (BleepingComputer) with evidentiary "media kits," accelerating public disclosure and pressure ahead of, or independent of, ransom payment.
- Willingness to pivot from public leak threats to private data sale when a sale is judged more commercially valuable than a leak (explicitly noted in the Novo Nordisk case, and as a general operating principle per Joe Shenouda's December 2025 analysis, which described an intent "to sell data to a single buyer" rather than leak publicly).

Claimed Data Volumes and Types

Victim (claim)	Claimed Volume	Claimed Data Types
Novo Nordisk	~1.3 TB / 700,000+ files	Source code, drug pipeline/research data, clinical trial records, employee/HCP/patient data, manufacturing details, internal AI models
Arup Group	~700GB GitHub + ~2TB cloud/DB	GitHub repos, Azure/AWS/S3 data, DB backups, BMS/ERP client data, code-signing certs, payment-gateway credentials, source code

Victim (claim)	Claimed Volume	Claimed Data Types
youX / DRIVE IQ	~300 GB	PII, government IDs, driver's licence numbers, loan application data, password hashes (MD5), broker employee credentials
Avnet	~1.3 TB (per later, expanded claim)	Financial models, customer data
Manchester Airports Group	~86 GB (claimed; unverified)	Customer profiles, booking/travel data, spend history, device/location data, ~200,000 upcoming-travel records

[CLAIM] All volume and content figures above are threat-actor claims. Independent verification exists only where explicitly noted (e.g., BleepingComputer's single validated MAG record; Novo Nordisk's own confirmation of a breach, though not of the specific 1.3TB/700,000-file figure).

Primary Focus Assessment

Assessed primarily as: Data theft and Extortion for Financial gain, with Publicity as an instrumental (not terminal) goal — media attention and leak-site "shaming" function to increase payment pressure, not as an end in themselves. No evidence of Disruption as a goal (no encryption/operational-impact incidents documented) and no evidence of Strategic impact motivation (no ideological, political, or state-aligned messaging identified in any reviewed source).

13. Infrastructure and Operational Security

[FACT] FulcrumSec maintains a dual clearnet/Tor leak-site presence (see Section 8) and uses encrypted, decentralised communication channels (Session, Tox) rather than centralised platforms that would be easier for researchers or law enforcement to subpoena or take down.

[ASSESSMENT] This infrastructure pattern (Tor leak site + decentralised messengers + a resilient clearnet mirror) is consistent with reasonably deliberate operational-security practice, though it is not unusual or especially sophisticated relative to other contemporary extortion groups (Moderate confidence; comparative, not a claim of unique sophistication).

[UNKNOWN] No reliable public evidence identified regarding FulcrumSec's hosting providers, ASN relationships, domain registration patterns/registrars, TLS certificate details, or DNS infrastructure patterns beyond the domain names themselves. No confirmed VPN/proxy infrastructure used by the group for operations (as opposed to for site hosting) was identified.

Given the absence of independently published network-infrastructure indicators (IP ranges, hosting-provider identification, certificate fingerprints), defenders should not expect to build network-layer detections against FulcrumSec-specific infrastructure at this time; detection efforts are better focused on the identity/credential/cloud-configuration behaviours documented in Sections 6–11.

14. Detection Opportunities — FulcrumSec Detection Matrix

Behavioural detections derived only from documented FulcrumSec activity, prioritised by how directly they map to the group's confirmed/claimed techniques.

Behaviour	Detection Source / Telemetry	Detection Logic (Summary)	Priority	Confidence
Secrets embedded in public-facing JavaScript/mobile bundles	External attack-surface scanning; CI/CD build artifacts; DevSecOps tooling	Automated regex/entropy scanning of deployed front-end code for API-key/token patterns (AWS, Azure, GitHub PAT, SaaS vendor key formats such as Iterable/Marketo).	Critical	High
Bulk repository cloning from a single token/identity	GitHub / Azure DevOps audit logs	Alert on a single PAT/service identity cloning an unusual number of repositories, or repositories outside its normal scope, in a short window.	High	High
Anomalous cloud-storage enumeration or download volume	AWS CloudTrail, Azure Activity Log, GCP Audit Log, MongoDB Atlas access logs	Alert on API calls enumerating buckets/collections followed by high-volume GetObject/export operations from a source not matching normal application behaviour.	High	High
Use of a single credential/token across multiple distinct cloud platforms in a short period	Cloud IAM / CIEM tooling	Correlate identity usage across AWS/Azure/GCP/SaaS to flag a token or service account suddenly active on platforms it has not previously accessed.	High	Moderate-High
Exploitation attempts against CVE-2025-55182 ("React2Shell") on React front-ends	WAF, application logs, EDR on web servers	Signature/behavioural detection for the published exploitation pattern; verify patch status across all internet-facing React deployments.	Critical (if unpatched)	High
Publicly exposed / open cloud storage or staging systems	Cloud security posture management (CSPM); external ASM scanning	Continuous scanning for public-read/write buckets, unauthenticated database endpoints (e.g., open MongoDB), and forgotten dev/staging subdomains.	High	High
JWT/authentication-bypass exposure ("alg:none")	Application security testing; API gateway logs	Verify all JWT validation logic explicitly rejects unsigned/`alg:none` tokens; log and alert on authentication events using unexpected algorithms.	Moderate	Moderate
Sustained low-and-slow outbound data	DLP; egress monitoring; UEBA	Baseline expected data-egress volumes per service identity; alert	Moderate	Moderate (based on

Behaviour	Detection Source / Telemetry	Detection Logic (Summary)	Priority	Confidence
transfer from a service account		on sustained anomalous transfer even below single-event volume thresholds (FulcrumSec has claimed deliberate throttling to evade volume-based detection in at least one case).		threat-actor claim)
Threat-actor pre-extortion contact with researchers/media referencing your organisation	OSINT/brand monitoring; breach-notification services	Monitor named breach-tracking outlets (BleepingComputer, DataBreaches.net), vx-underground, and social media for early mentions of your organisation alongside "FulcrumSec."	Moderate	High (behavioural pattern well documented)

15. IOC Catalogue

IOC Availability Notice

FulcrumSec is a young, cloud/identity-focused actor with no publicly documented malware, so the conventional IOC categories most useful for classic malware-based actors (file hashes, mutexes, registry keys, C2 IPs) have essentially no public evidence base for this group. This dossier lists only IOC-adjacent items with direct public sourcing and does not fabricate hashes, IPs, or other indicators to fill the requested template.

Indicator	Type	Context	Source	Status
fulcrumsec[.]net	Domain	Group-operated clearnet leak/shaming site ("Index of /Shame").	Multiple OSINT trackers; direct observation	Attributed / Active (as of source dates)
fulcrumsec[.]vg	Domain	Group-operated campaign/victim listing page.	OSINT tracking	Attributed / Active (as of source dates)
FulcrumSec Tor (.onion) leak site (specific address not independently re-verified for this dossier)	URL (Tor)	Full stolen-data publication for non-paying victims.	Referenced across multiple victim disclosures (e.g., Arup, Avnet)	Attributed / status uncertain
Threatspians@fulcrumsec[.]net (exact address as reported)	E-mail address	Victim/researcher/media contact channel.	Dataminr Cyber Intel Brief, March 2026	Attributed / Historical (confirm current validity before use)
CVE-2025-55182 ("React2Shell")	Vulnerability ID	Sole publicly documented CVE used by the group; unauthenticated RCE on unpatched React front-ends.	MOXFIVE, June 2026	Attributed / Confirmed vulnerability, ongoing relevance
"FulcrumSec" / "The Threat Thespians" (brand names/strings)	Actor name / string	Useful as a threat-intelligence monitoring and log-scraping search term per Joe Shenouda's recommendation.	shenouda.nl, December 2025	Attributed

[UNKNOWN] No file hashes, malware samples, mutexes, registry keys, TLS certificate fingerprints, user agents, or cryptocurrency wallet addresses attributable to FulcrumSec were identified in this research. No reliable public evidence identified for these categories. Do not treat their absence here as confirmation they do not exist — treat it as confirmation that no corroborated public source was located as of 2 September 2026.

16. Defensive Strategy

Identity

- Enforce MFA universally, including for service/API accounts and CI/CD systems wherever the platform supports it.
- Move from long-lived personal access tokens to short-lived, scoped, auto-rotating credentials for all repository, cloud, and third-party SaaS integrations.
- Deploy identity monitoring/CIEM to flag a single credential accessing multiple unrelated cloud platforms or an unusual volume of resources.
- Apply conditional access policies restricting service-account and API-key use to expected source ranges and usage patterns.

Endpoint

- FulcrumSec's documented activity does not rely on endpoint malware; endpoint controls (EDR, application control) have limited direct relevance to this specific actor but remain part of baseline hygiene against unrelated threats.
- Where public-facing application servers (e.g., React front-ends) run on monitored endpoints, ensure EDR/HIDS coverage extends to those hosts to catch post-exploitation activity following CVE-2025-55182-style RCE.

Network

- Egress filtering and anomaly detection for cloud-service and SaaS-API traffic volumes, since FulcrumSec's exfiltration occurs over standard web/API channels rather than custom C2 protocols.
- Segment cloud accounts and environments by business unit/entity to contain the blast radius of a single compromised credential (see Section 11).

Email

- No FulcrumSec-specific phishing campaign has been documented; nonetheless, standard DMARC/DKIM/SPF enforcement and attachment/URL filtering remain relevant given exfiltrated contact data (e.g., from the MAG claim) is likely to fuel copy-cat phishing/smishing campaigns referencing real booking or account details.
- Pre-position customer-facing guidance (as MAG did) proactively warning customers that the organisation will never request payment or account details via unsolicited contact.

Cloud

- This is the primary control domain for FulcrumSec risk reduction: enforce least-privilege IAM policies, enable and centrally monitor cloud audit logging (CloudTrail/Azure Activity Log/GCP Audit Log) with retention sufficient for post-incident investigation, and run continuous CSPM scanning for public storage exposure.
- Extend the same rigor to third-party SaaS platforms integrated via API (marketing/CX platforms such as Iterable, Marketo, Qualtrics; data platforms such as Databricks, MongoDB Atlas) — these sit outside traditional "cloud security" scope but were the specific access point in the MAG claim and multiple other cases.

Data

- Deploy DLP and sensitive-data discovery/classification across cloud storage and SaaS platforms, not only on-premises file shares.
- Ensure backups and any archived customer/marketing exports are subject to the same access controls and monitoring as production data — FulcrumSec has specifically targeted "export"-style consolidated datasets (e.g., the claimed 21.5GB "Manchester customer export").

Vulnerability Management

- Prioritise patching of CVE-2025-55182 ("React2Shell") on any internet-facing React application as an immediate action — this is the group's only publicly documented CVE and has been observed unpatched for months in at least one victim environment.
- Prioritise discovery and decommissioning/protection of "forgotten" development and staging subdomains, which have been the specific entry point in multiple cases (Novo Nordisk, Arup).

Incident Response

- Build an IR playbook assuming credential-based, malware-free intrusion: focus initial triage on cloud/SaaS audit logs and identity activity rather than endpoint forensics alone.
- Prepare communications processes for the possibility that an attacker will contact media/researchers directly and rapidly (as FulcrumSec did with BleepingComputer within days of a victim's own disclosure).
- On suspected compromise, immediately rotate all credentials/tokens reachable from the initially exposed one (assume repository/token compromise cascades to other systems, per the Novo Nordisk and Arup pattern) rather than rotating only the initially identified credential.

17. OT / ICS / Critical Infrastructure Implications

Explicit Scoping Statement

No reliable public evidence identifies FulcrumSec as having attacked, targeted, or demonstrated capability against OT, ICS, SCADA, building management systems, airport operational technology, industrial control systems, safety systems, physical security/access control, CCTV, communications infrastructure, baggage systems, airside systems, or energy systems — including in the claimed MAG intrusion. MAG's own statement and all independent reporting reviewed are consistent in stating that airport operations, passenger safety, and aviation security were unaffected.

Observed OT Activity

[UNKNOWN] None. No reliable public evidence identified of any FulcrumSec activity touching OT/ICS/BMS/safety/physical-security systems in any documented case (MAG or otherwise).

Indirect OT Exposure

[CLAIM] The one partial exception across all documented FulcrumSec cases is the Arup Group claim (May 2026), in which the group claimed to have obtained "Neuron BMS client databases" — i.e., data related to a building-management-system product used by Arup's clients — as part of a broader claimed data haul. This is a claim about data pertaining to a BMS product/its clients, not evidence of an intrusion into, or manipulation of, an operational BMS deployment itself. Threat-actor claim — independently unverified, and even if accurate, describes data exposure rather than OT compromise.

Plausible Attack Paths (Analytical, Not Observed)

The following are Shieldworkz analytical constructs describing theoretically plausible pathways by which an actor with FulcrumSec's demonstrated technique set (credential/API-key hunting, cloud misconfiguration exploitation) could indirectly create OT-relevant risk. These are not claims that FulcrumSec has done or intends to do this — they are offered to inform proactive OT defensive planning, per the research brief's request to distinguish plausible paths from observed activity.

- ASSESSMENT (speculative/plausible, not observed): A cloud-hosted BMS/asset-management platform (of the kind referenced in the Arup/Neuron BMS claim) that stores client site configuration, network diagrams, or remote-access credentials could, if compromised via FulcrumSec's documented credential-hunting technique, provide reconnaissance value to any actor seeking to plan a subsequent OT-directed attack — but this would require a separate actor or a deliberate pivot beyond anything FulcrumSec has been observed to do.
- ASSESSMENT (speculative/plausible, not observed): Airport or critical-infrastructure operators commonly use shared cloud/SaaS marketing and customer-engagement platforms (the class of system implicated in the MAG claim) that are organisationally and technically separate from OT networks; the plausible risk path is reputational/regulatory and

customer-fraud, not a technical bridge into OT, absent evidence of shared credentials or flat network architecture linking the two domains (none identified in the MAG case).

Unsupported Speculation (Explicitly Rejected)

[UNKNOWN] This dossier explicitly does not claim, and finds no basis to claim, that FulcrumSec has the intent, capability, or documented history to conduct OT/ICS-disruptive attacks, or that the MAG incident had any bearing on airport safety or operational systems. Any such claim circulating elsewhere should be treated as unsupported absent new evidence.

OT Defensive Considerations (Evidence-Based, Precautionary)

Notwithstanding the absence of observed OT targeting, the following controls are recommended for OT/critical-infrastructure security teams specifically because FulcrumSec's demonstrated technique set (credential exposure, cloud misconfiguration, third-party SaaS API abuse) is directly transferable risk wherever OT-adjacent data or systems share those same weaknesses:

5. Maintain strict logical and credential separation between OT/ICS networks and IT/cloud/marketing systems — verify no shared service accounts, API keys, or identity providers span both domains.
6. Extend secrets-scanning and credential-hygiene programmes (Section 10/16) to any cloud-hosted OT-adjacent platforms: BMS management portals, remote-monitoring dashboards, vendor support/access tools, and asset-management systems — the class of system referenced in the Arup/Neuron BMS claim.
7. Treat any third-party engineering, BMS, or facilities-management vendor with the same third-party risk scrutiny applied to core IT vendors, since vendor-side data exposure (as claimed against Arup) could reveal OT network topology, credentials, or configuration data useful to a different, OT-focused actor even where FulcrumSec itself has no OT intent.
8. Ensure OT/ICS incident-response plans include a defined communication and verification path with IT/cloud security teams so that an IT-side credential-exposure incident (of the kind FulcrumSec specializes in) can be rapidly assessed for any OT-adjacent blast radius, even when — as in the MAG case — the initial indications point away from OT impact.

18. Sector-Specific Risk

Sector	Why It Matters	Likely Attack Surface	Critical Dependencies	Recommended Controls
Airports / Aviation	Directly implicated via the claimed MAG intrusion; high customer-data volume, high public/media visibility, mandatory regulatory reporting (ICO/NCSC).	Customer-facing booking, parking, lounge, Wi-Fi, and CX/marketing SaaS integrations (per the MAG claim).	Third-party marketing/CRM platforms, customer-data warehouses, booking systems.	API credential hygiene for all customer-engagement SaaS; segmentation from OT/operational systems (validated as effective in the MAG case); customer phishing-awareness readiness.
Airlines	Shares the aviation customer-data profile and regulatory exposure of airports; not a documented FulcrumSec victim to date.	Booking/loyalty platforms, customer-data warehouses, partner API integrations.	GDS/booking-system integrations; loyalty-program data stores.	Same credential/API hygiene priorities as airports; monitor for exposed keys in web/mobile app code.
Ports / Rail	No documented FulcrumSec activity; included per audience requirement as adjacent critical-transport infrastructure.	Assessed by analogy to airports: customer-facing digital services, not documented OT-relevant surface.	Ticketing/booking platforms; freight/logistics data systems.	Apply the same customer-data-platform hygiene reviewed in Section 16, pending any sector-specific evidence.
Energy	No documented FulcrumSec activity or claim.	No reliable public evidence identified.	N/A — no evidence base.	General cloud-credential hygiene per Section 16 as baseline practice; no FulcrumSec-specific control indicated.
Manufacturing / Distribution	Confirmed sector via Avnet (electronics distribution).	Sales/CRM tools, cloud storage, internal repositories.	ERP/CRM integrations; supplier/customer data platforms.	Secrets-scanning across sales-tool integrations; API key rotation for third-party sales platforms.
Government	No documented FulcrumSec victim.	No reliable public evidence identified.	N/A	No FulcrumSec-specific control indicated; general public-sector cloud-security baselines apply.
Healthcare / Pharmaceuticals	Confirmed sector via Novo Nordisk; claimed victim Lena Health.	Clinical/research cloud platforms, dev/sandbox subdomains, source-code repositories.	Clinical trial data systems, HCP/patient data stores, AI/ML model repositories.	Rigorous separation and access control for dev/sandbox environments; JWT/authentication hardening (reject unsigned tokens); repository secrets-scanning.

Sector	Why It Matters	Likely Attack Surface	Critical Dependencies	Recommended Controls
Financial Services / FinTech	Confirmed/claimed via youX/DRIVE IQ and LexisNexis Risk Solutions.	Cloud databases (MongoDB, Redshift), customer-facing web applications.	Loan/broker data platforms; risk-data aggregation services.	Patch internet-facing web frameworks against known CVEs (e.g., CVE-2025-55182); enforce least-privilege container/IAM permissions; secrets-manager hygiene.
Technology	General sector inclusion per multiple vendor trackers; broad claimed victim base.	Source-code repositories, CI/CD pipelines, cloud infrastructure.	DevOps toolchains; customer data platforms.	CI/CD secrets hygiene; repository access auditing; least-privilege service accounts.
Other Critical Infrastructure (general)	No FulcrumSec-specific evidence beyond MAG; included for completeness given audience.	Assessed by analogy: customer-facing and third-party SaaS integrations are the most plausible exposure class based on the group's documented pattern.	Third-party vendor/SaaS integrations generally.	Apply Section 16 defensive strategy as a sector-agnostic baseline; prioritise based on actual cloud/SaaS footprint.

19. Threat Hunting Playbook

Hypotheses are built only from verified/strongly supported FulcrumSec behaviour (Sections 6–11).

Hypothesis 1 — Exposed secrets in public-facing code

- Hypothesis: An API key, access token, or credential for a cloud or SaaS platform is exposed in publicly served JavaScript, mobile app binaries, or a public repository.
- Data sources: External attack-surface scan results; source-control repository content and history; deployed web/app bundle artifacts.
- Investigative approach: Run automated secret-pattern/entropy scanning against all public-facing JS bundles and repositories; for any hit, verify whether the exposed credential is live and what scope/permissions it carries.
- Positive finding: A live, in-scope credential (e.g., a marketing-platform API key, cloud storage credential, or PAT) reachable from publicly served code.
- Follow-up investigation: Determine whether the credential has been used from any unexpected source IP/location/time (via the issuing platform's audit log) prior to remediation.
- Containment: Immediately revoke/rotate the credential; review the issuing platform's access logs for anomalous prior use; assume compromise until logs prove otherwise.

Hypothesis 2 — Bulk repository or data-platform access from a single credential

- Hypothesis: A single token or service identity has accessed or cloned an unusual volume/scope of repositories or cloud-storage objects.
- Data sources: GitHub/Azure DevOps audit logs; AWS CloudTrail; Azure Activity Log; GCP Audit Log; Databricks/MongoDB Atlas access logs.
- Query logic (Splunk, illustrative): `index=github_audit action=git.clone OR action=repo.download | stats count dc(repo) as repos by actor_id | where repos > <baseline_threshold>`
- Query logic (Microsoft Sentinel, illustrative, KQL): `AuditLogs | where OperationName in ("Clone repository","Download repository") | summarize RepoCount=dcount(TargetResources) by InitiatedBy, bin(TimeGenerated, 1h) | where RepoCount > <baseline_threshold>`
- Query logic (Elastic, illustrative, EQL/DSL concept): aggregate `event.action:"repo_clone"` by `user.id` over a rolling window, alert where distinct-repo count exceeds a per-identity historical baseline.
- Positive finding: A single identity accessing significantly more repositories/objects than its historical baseline, especially outside normal business hours or from an unrecognised network context.
- Follow-up: Identify what the accessed repositories/objects contain (search for embedded secrets) and whether any of those secrets have since been used elsewhere.
- Containment: Suspend the identity, rotate all credentials it could reach, and force re-authentication for downstream systems.

Hypothesis 3 — Unpatched CVE-2025-55182 ("React2Shell") exposure

- Hypothesis: An internet-facing React application in the environment remains unpatched against CVE-2025-55182.
- Data sources: Vulnerability-management/asset-inventory data; WAF and web-server access logs.
- Investigative approach: Cross-reference the vulnerability scanner's CVE database against all internet-facing web application inventories; review WAF/access logs for known exploitation request patterns for this CVE.
- Positive finding: An unpatched, internet-facing instance, or log evidence of exploitation attempts/successful requests matching the published pattern.
- Follow-up: Full incident-response engagement, including a check for webshell or unauthorised code artifacts on the affected host.
- Containment: Emergency patch or temporary WAF virtual-patch rule; isolate the host pending investigation if exploitation indicators are found.

Hypothesis 4 — Sustained low-volume anomalous egress from a service identity

- Hypothesis: A service account or API identity is transferring data at a sustained, low, but abnormal rate consistent with deliberate detection evasion (per FulcrumSec's claimed throttling behaviour).
- Data sources: Cloud storage access/download logs; DLP; network egress monitoring; UEBA baselines.
- Query logic (Sentinel/KQL, illustrative): `StorageBlobLogs | summarize TotalBytes=sum(ResponseBodySize) by AccountName, bin(TimeGenerated, 1d) | where TotalBytes > baseline_daily_avg * <multiplier> and TotalBytes < single_event_alert_threshold` — designed to catch sustained-but-subthreshold transfer that single-event volume alerts would miss.
- Positive finding: A multi-day pattern of elevated-but-individually-modest data egress from one identity, inconsistent with its normal application function.
- Follow-up: Identify what data was transferred and to where (destination IP/service); correlate with any concurrent credential-exposure findings from Hypotheses 1–2.
- Containment: Suspend the identity pending investigation; engage IR for scoping.

All query logic above is illustrative and framework-appropriate, not vendor-certified syntax; adapt field names and thresholds to your environment before deployment.

20. Incident Response Playbook

A FulcrumSec-specific response sequence, adapted from the group's documented credential/cloud-based TTPs rather than a generic ransomware playbook.

Stage	Key Actions
Detection	Trigger from any Section 14/19 detection (exposed secret found, anomalous repo/cloud access, CVE-2025-55182 exploitation indicator, sustained anomalous egress, or external notification such as media/researcher contact referencing your organisation and "FulcrumSec.")
Triage	Establish scope of the initially identified exposure: which credential/system, what permissions it carried, and what it could reach. Determine whether this matches FulcrumSec's documented pattern (credential exposure → repository/cloud access → bulk data collection) to calibrate urgency.
Scoping	Given FulcrumSec's documented "credential cascade" pattern, scope outward from the initial exposure: identify every credential/secret reachable from the compromised one (e.g., other tokens stored in the same repository), not only the originally identified item.
Containment — Credential Containment	Revoke/rotate the exposed credential immediately, then revoke/rotate every credential it could plausibly have exposed (repositories cloned, secrets managers accessed). Do not assume a single rotation is sufficient.
Containment — Network Isolation	For any case involving CVE-2025-55182 or similar public-facing exploitation, isolate the affected host/application pending forensic review. Network isolation is of secondary relevance for pure credential/cloud-API cases (there may be no compromised host to isolate).
Evidence Preservation	Preserve cloud/SaaS audit logs (many retain limited history by default — extend/export immediately), repository access logs, and any communications received from an actor claiming responsibility, including e-mail headers and data samples provided as "proof."
Eradication	Confirm no residual access remains via any credential in the exposure chain; verify no unauthorised repository forks, cloud-storage replication jobs, or scheduled exports were established during the dwell period.
Recovery	Restore normal service using newly issued, scoped credentials; re-enable any suspended integrations only after confirming least-privilege scoping.
Threat Hunting	Execute Section 19 hypotheses across the broader environment (not just the initially affected system) to rule out parallel or preceding exposure, given FulcrumSec's documented pattern of extended, multi-month dwell time in some cases.
Monitoring	Establish ongoing monitoring for the specific credential/system class involved (per Section 14) and for any leak-site or media mention of your organisation.
Lessons Learned	Feed findings into secrets-scanning coverage, credential-rotation policy, and vendor/API-integration review; assess whether the exposure class (client-side JS, repository, or open storage) exists elsewhere in the environment.

Decision Points

- Suspected credential compromise: Treat as an active exposure until logs prove otherwise; rotate first, investigate scope second — the group's documented pattern rewards delay.
- Suspected data exfiltration (no ransomware note or encryption observed): Do not wait for an extortion contact to confirm exfiltration; FulcrumSec's model produces no encryption artifact, so absence of ransomware indicators does not indicate absence of a breach.
- Suspected third-party/SaaS-vendor compromise (e.g., a marketing/CX platform): Engage the vendor's own incident-response process in parallel with internal investigation; per the MAG claim, the initial exposure point may sit at the vendor-integration layer rather than in owned infrastructure.
- Suspected privileged-account or repository compromise: Assume credential cascade (Section 10/11) and scope outward before declaring containment complete.
- Suspected OT impact: No FulcrumSec case has shown OT impact; nonetheless, verify segmentation integrity as a precaution per Section 17, and do not assume separation without positive confirmation during the incident.

21. Early Warning Indicators

Actor-Specific Indicators (FulcrumSec)

- Discovery of a live, exploitable API key or token in your own or a partner's publicly served JavaScript, mobile app, or public repository — this is the group's primary and best-evidenced entry point.
- An internet-facing React application confirmed unpatched against CVE-2025-55182.
- Discovery of a publicly accessible cloud storage bucket, staging server, or database that would qualify for the group's "Index of /Shame" targeting criteria.
- Unexplained, sustained (multi-week to multi-month) low-level anomalous access to source-code repositories or cloud storage from a single credential — consistent with the group's documented extended-dwell-time pattern.
- Early mention of your organisation alongside "FulcrumSec" or "The Threat Thespians" on breach-tracking sites, vx-underground, DataBreaches.net, or in direct contact from BleepingComputer or similar outlets seeking comment.

Generic Ransomware/Extortion Indicators (Not FulcrumSec-Specific)

- Unusual authentication patterns (impossible-travel logins, off-hours privileged access) — a general indicator applicable to many actors, not confirmed as a FulcrumSec-specific pattern.
- Abnormal VPN activity or use of unauthorised remote-access tools — no evidence FulcrumSec uses these vectors; listed here only because the research brief requests general early-warning coverage, and flagged explicitly as not actor-specific.
- Anomalous privilege escalation via classic Windows/AD techniques — not documented for FulcrumSec; a generic indicator only.
- Abnormal outbound traffic and suspicious DNS behaviour in the classic C2-beaconing sense — FulcrumSec's documented exfiltration uses standard cloud/SaaS channels rather than covert C2, so classic beaconing detections are of limited relevance to this specific actor.
- Unusual cloud activity and leak-site activity generally — these overlap with the actor-specific indicators above and remain broadly relevant.

This separation matters operationally: SOC teams should not down-rank or dismiss a FulcrumSec-relevant finding (e.g., an exposed API key) simply because it does not resemble the "generic ransomware precursor" indicators (VPN abuse, AD reconnaissance) that dominate many playbooks — FulcrumSec's risk profile is materially different.

22. Unknowns and Intelligence Gaps

What We Still Do Not Know

Intelligence Gap	Why It Matters	What Evidence Would Resolve It	How Defenders Can Monitor
Exact organisational structure and size	Determines whether FulcrumSec is a small crew or a larger, possibly multi-cell operation, affecting expected operational tempo and resilience to disruption.	Law-enforcement investigation results; leaked internal communications; researcher infiltration of the group's communication channels.	Track researcher/journalist reporting (BleepingComputer, DataBreaches.net, SecurityWeek) for any future law-enforcement action or leaked internal detail.
True operator identities	Needed for legal/law-enforcement action and for assessing any individual operator's prior history/capability.	Indictment, arrest, or confirmed doxxing by a credible investigator.	Monitor law-enforcement press releases (NCA, FBI, Europol) and established CTI vendor reporting for named individuals.
Funding model and financial scale	Affects assessment of the group's sustainability and incentive to continue/escalate operations.	Blockchain/financial-forensics tracing of any confirmed ransom payments; law-enforcement financial investigation results.	Track for any confirmed ransom payment reporting tied to a specific, traceable wallet.
Geographic location of operators	Relevant to law-enforcement jurisdiction and to any (currently unsupported) state-nexus assessment.	Law-enforcement attribution; infrastructure/registration forensics not currently public.	Monitor for any government advisory naming a country of origin.
Relationships with other groups (e.g., "TheUSERS007")	Determines whether FulcrumSec is a standalone brand or part of a looser criminal ecosystem/affiliate structure.	Confirmed shared infrastructure, personnel overlap, or coordinated-campaign evidence from independent researchers.	Track future joint or overlapping victim claims and any researcher commentary on infrastructure overlap.
Full victim set	Current victim counts (21–25+) are all self-reported by the actor via its leak site; true scope, including unreported/paid victims, is unknown.	Victim self-disclosure, regulatory filings, or independent breach-notification datasets.	Monitor regulator breach-notification registers (e.g., ICO, state AG breach registers) for unattributed incidents matching FulcrumSec's TTP profile.
Complete infrastructure footprint	Only two domains, a Tor site, and two messaging identifiers are documented; hosting, ASN, and historical domain infrastructure are not public.	Independent infrastructure-mapping research (passive DNS, certificate transparency analysis) not identified in current public reporting.	Monitor CTI vendor blogs for any future infrastructure-mapping publication.
Full toolset	No malware has been documented; it is unknown whether this reflects a genuinely malware-free	A forensic incident-response engagement with public/shareable findings from a FulcrumSec victim.	Track CTI vendor blogs (Sysdig, MOXFIVE, Mandiant, CrowdStrike, Microsoft) for

Intelligence Gap	Why It Matters	What Evidence Would Resolve It	How Defenders Can Monitor
	operating model or simply a lack of forensic access by researchers to date.		future FulcrumSec-specific technical reporting.
True initial-access preferences vs. documented cases	Three "doors" are documented, but this may reflect what has been publicly disclosed rather than the group's full repertoire.	Additional victim forensic disclosures.	Continue aggregating new claimed/confirmed victim technical detail as it is published.
Operational security practices	Affects law-enforcement disruption prospects and the durability of the group's leak-site/communication infrastructure.	Takedown attempts and their success/failure; researcher infrastructure analysis.	Monitor for leak-site downtime/migration patterns as an indirect signal.
Actual capabilities vs. claimed capabilities	The group has an incentive to inflate claimed data volumes and sophistication for extortion leverage (illustrated by the unresolved gap between MAG's disclosure and FulcrumSec's 86GB claim).	Independent forensic verification of claimed data volumes in any case.	Track for any future independent (non-victim, non-actor) technical verification of claimed breach scope.
Extent of any OT capability	Currently zero evidence of intent or capability; this could change, and CISOs should not become complacent based on a point-in-time absence of evidence.	Any future claimed or confirmed OT-adjacent intrusion.	Monitor future leak-site postings and victim disclosures for any OT/ICS/BMS-operational (not just BMS-data) claim.
Future targeting	Cannot be forecast with confidence; the group's only consistent pattern is technique-driven opportunism, not sector loyalty.	N/A — inherently forward-looking and unknowable in advance.	Maintain the Section 14/19 detections as standing controls rather than reactive, victim-specific ones, since the next target is unpredictable by sector.

23. Myth vs Fact

Claim in Circulation	Classification	Evidence / Basis
"FulcrumSec deploys ransomware and encrypts victim systems."	False / Misleading	No ransomware binary or encryption behaviour has been documented in any reviewed case; MOXFIVE and Sysdig both explicitly confirm the group is data-theft-only ("steal and squeeze").
"FulcrumSec caused the Manchester Airports Group breach."	Unverified (Probable per threat-actor claim, but not independently confirmed)	FulcrumSec's own claim to BleepingComputer, partially corroborated by one validated data sample; MAG has not confirmed the attacker's identity. See Section 4.6.
"The MAG breach compromised passenger safety or airport operations."	False	MAG, NCSC-adjacent reporting, and all independent outlets reviewed are consistent: no operational or safety-system impact occurred.
"FulcrumSec used a zero-day exploit against MAG or other victims."	Unverified / Probably False	The only CVE publicly and reliably tied to FulcrumSec (CVE-2025-55182) is a known, patched vulnerability, not a zero-day. No zero-day use has been documented for MAG or any other case.
"FulcrumSec is affiliated with LAPSUS$/Karakurt/ShinyHunters."	Unverified	Researchers have drawn operational/stylistic comparisons (similar identity-based, malware-free extortion model) as analytical context, not as evidence of shared membership, coordination, or lineage. No confirmed affiliation exists.
"FulcrumSec never targets healthcare organisations."	False / Misleading	The group reportedly stated this intent in early victim communications (per MOXFIVE), but healthcare victims (Lena Health) subsequently appeared on its leak site — a documented inconsistency between stated principle and observed conduct.
"FulcrumSec is a nation-state-sponsored actor."	Unverified / No Evidence	No researcher or government source reviewed attributes state sponsorship; all corroborated reporting characterises the group as financially motivated, criminal, and non-state.
"The Manchester Airports Group breach exposed banking or payment card data."	False	Confirmed by MAG's own statement, corroborated by BleepingComputer's sample review, that no such data was held in the affected system.
"FulcrumSec exfiltrated 86GB of MAG data."	Unverified (Claim, single sample corroborated)	This is the threat actor's figure; BleepingComputer validated one record's authenticity but explicitly could not verify total volume.

24. Intelligence Confidence Assessment

Overall confidence ratings for the dossier's major findings, using High / Medium / Low confidence per standard analytic tradecraft. Confidence reflects the strength, independence, and consistency of sourcing — not the severity of the underlying risk.

Finding	Confidence	Rationale
FulcrumSec is a financially motivated, malware-free, cloud/credential-focused data-extortion group active since ~September 2025.	High	Convergent, independent reporting from Sysdig, MOXFIVE, Dataminr, Mallory.ai, and multiple victim-adjacent disclosures, with consistent technical detail across more than a dozen cases.
FulcrumSec's primary initial-access vectors are exposed credentials/API keys, one specific known CVE (CVE-2025-55182), and open cloud storage.	High	Explicitly documented and cross-referenced by MOXFIVE and Sysdig across multiple distinct victims with consistent technical pattern.
FulcrumSec was responsible for the August 2026 Manchester Airports Group data breach.	Medium	Actor self-claim to a reputable outlet, one independently validated data sample, and consistency with the group's known TTPs — but no victim, forensic, or law-enforcement confirmation.
The claimed ~86GB / ~200,000 upcoming-travel-record MAG dataset figures are accurate.	Low	Entirely threat-actor sourced; independent verification exists only for one individual record, not for aggregate volume.
No OT/ICS/aviation-safety systems were affected in the MAG incident.	High	Consistent, unchallenged statements from MAG and multiple independent outlets across a week-plus reporting period, with no contrary reporting identified.
FulcrumSec has no relationship with, or is not affiliated with, any specific other named threat actor or nation-state.	Low-Medium (absence of evidence, not evidence of absence)	No positive evidence of affiliation was found, but the group's youth and limited researcher access to its internals mean a currently undetected affiliation cannot be excluded.

Where this dossier could not assign meaningful confidence (e.g., organisational structure, funding, operator identity — see Section 22), it states the gap explicitly rather than assigning a placeholder confidence level to speculation.

25. CISO Executive Takeaways

9. FulcrumSec is not a ransomware threat in the classic sense — it never encrypts. The actual threat is a rapid, quiet, malware-free data-theft-and-extortion operation that most encryption-focused detection and resilience investments will not catch.
10. What makes FulcrumSec different is its almost total reliance on credential and configuration hygiene failures — exposed API keys, unpatched web frameworks, and open cloud storage — rather than novel malware or zero-day exploits. This makes the group unusually preventable with disciplined, unglamorous controls.
11. The single highest-leverage control is secrets management: scanning for and eliminating hardcoded/exposed credentials in client-side code, mobile apps, and source repositories has more documented relevance to this actor than almost any other control category.
12. Patch CVE-2025-55182 ("React2Shell") on every internet-facing React deployment now — it is the only CVE publicly and reliably tied to this group, and it has been observed unpatched for months in at least one victim.
13. Third-party marketing/customer-engagement SaaS platforms (the class implicated in the MAG claim) deserve the same credential-hygiene rigor as core infrastructure — they are not "low-risk" simply because they sit outside traditional IT/OT security scope.
14. The Manchester Airports Group incident is not established as a confirmed FulcrumSec attack — it is a plausible, partially corroborated claim. Treat the attribution as Moderate confidence, not fact, when briefing your own board or regulators, and expect the picture to evolve.
15. The clearest confirmed lesson from the MAG incident is a positive one: segmentation between customer-facing digital services and safety/operational systems held. This is validation, not an excuse to under-invest in it going forward.
16. If you detect activity matching FulcrumSec's pattern (a live exposed credential, anomalous bulk repository/cloud access, sustained low-level anomalous egress), assume a credential cascade and rotate outward from the initial finding — do not treat the first exposed secret as the full scope.
17. FulcrumSec has shown it will proactively contact media and researchers to increase extortion pressure, often within days of a victim's own disclosure. Incident communications plans should assume this possibility, not treat it as a low-likelihood edge case.
18. There is no credible evidence FulcrumSec has attacked or intends to attack OT/ICS/safety systems, including in the MAG case. Do not let MAG headlines drive disproportionate reallocation of OT security budget away from evidenced OT-specific threats.
19. FulcrumSec's targeting is opportunistic and technique-driven, not sector-loyal. No sector should assume it is "not a target" — the deciding factor has consistently been exploitable weakness, not industry.
20. Victim disclosures and attacker claims have diverged meaningfully in scope in this group's history (MAG, Novo Nordisk); plan customer/regulator communications around the possibility that the true exposure could exceed the initial official disclosure.
21. A large share of what is publicly known about FulcrumSec comes from the group's own leak-site claims and self-reported "campaign" framing. Treat victim counts, data volumes,

and technical narratives from the actor as marketing for its extortion operation, not as neutral fact, until independently corroborated.

22. Significant intelligence gaps remain (operator identity, true scale, funding, affiliations — Section 22). Build detection and defence around the group's demonstrated technique set, not around assumptions about its size or backing that current evidence cannot support.

26. CISO Action Checklist

Immediate — 0–7 Days

- Run an emergency scan of all public-facing JavaScript, mobile app bundles, and public repositories for embedded API keys/credentials for cloud and SaaS platforms (especially customer-engagement/marketing platforms such as Iterable, Marketo, Qualtrics).
- Confirm patch status of every internet-facing React application against CVE-2025-55182; apply emergency WAF virtual-patching for anything not yet remediated.
- Inventory and review permissions/scope for every third-party SaaS integration that touches customer data, prioritising marketing/CX platforms.
- Brief the incident-communications team on the possibility of proactive attacker outreach to media/researchers, using the MAG/BleepingComputer sequence as a concrete example.

Near Term — 7–30 Days

- Deploy or expand automated secrets-scanning across CI/CD pipelines, Git history, and deployed front-end artifacts, with mandatory rotation workflows for any finding.
- Implement or tune cloud audit-log monitoring (CloudTrail/Azure Activity Log/GCP Audit Log, plus GitHub/Azure DevOps audit logs) for the anomalous bulk-access patterns described in Section 14.
- Conduct a cloud-storage exposure sweep (CSPM/external ASM) to identify and remediate any publicly accessible buckets, databases, or forgotten dev/staging subdomains.
- Review and, where feasible, move long-lived personal access tokens/API keys to short-lived, scoped, auto-rotating credentials for the highest-risk integrations identified above.

Medium Term — 30–90 Days

- Establish or formalise a third-party/vendor risk review specifically covering API-integrated SaaS platforms (not only traditional data-processing agreements), incorporating credential-hygiene attestations.
- Build the Section 19 threat-hunting hypotheses into a recurring hunt cadence rather than a one-time exercise.
- Run a tabletop exercise simulating a FulcrumSec-style incident: credential exposure discovered, no ransomware note, attacker contacts a journalist before the organisation has completed internal scoping.
- Extend cloud identity governance (CIEM) coverage to flag single-identity access spanning multiple unrelated platforms, per Section 11.

Strategic — 90+ Days

- Embed secrets-management and least-privilege cloud/SaaS credential architecture as a standing engineering requirement (secure-by-design), not a point-in-time remediation.

- Establish continuous, org-wide external attack-surface monitoring (not periodic/manual) given how quickly a forgotten subdomain or exposed key can become a multi-million-record breach.
- Maintain a standing OSINT/brand-monitoring watch for organisational mentions alongside FulcrumSec or other emerging data-extortion actors, feeding into early-warning processes (Section 21).
- Revisit this dossier's assessments periodically — FulcrumSec is an actively evolving, young threat actor, and confidence levels in Sections 4 and 24 in particular should be expected to change as the MAG incident and the group's broader activity develop.

27. Appendix

27.1 Glossary

Term	Definition
Double extortion	An extortion model combining data theft with a second pressure mechanism (typically encryption/operational disruption). FulcrumSec's model theft-plus-publish-threat resembles this in effect but omits the encryption component.
Steal-and-squeeze	FulcrumSec's self-described operating model: data theft without encryption, followed by extortion and, if unpaid, publication or private sale.
CVE	Common Vulnerabilities and Exposures — a standardised identifier for a publicly known software vulnerability.
Dwell time	The period between an attacker gaining initial access and being detected or disclosed.
CIEM	Cloud Infrastructure Entitlement Management — tooling for governing and monitoring cloud identity permissions.
CSPM	Cloud Security Posture Management — tooling for continuously identifying cloud misconfigurations.
PAT	Personal Access Token — a credential used to authenticate to a code-hosting or API platform, often as an alternative to a password.

27.2 Acronyms

Acronym	Meaning
MAG	Manchester Airports Group
ICO	Information Commissioner's Office (UK data protection regulator)
NCSC	National Cyber Security Centre (UK)
GDPR	General Data Protection Regulation
TTP	Tactics, Techniques and Procedures
IOC	Indicator of Compromise
DLP	Data Loss Prevention
RCE	Remote Code Execution
JWT	JSON Web Token

Acronym	Meaning
BMS	Building Management System
OT/ICS	Operational Technology / Industrial Control Systems

27.3 Confidence Methodology

This dossier uses a four-tier evidentiary tagging system (FACT / ASSESSMENT / CLAIM / UNKNOWN) applied at the level of individual statements, plus a standard analytic-tradecraft confidence scale (High / Medium / Low, and — for the MAG attribution specifically — Confirmed / High / Moderate / Low / Unconfirmed) applied to overall findings. A statement is tagged FACT only where corroborated by a primary source (victim/regulator/official statement) or independently verified technical detail. A statement is tagged CLAIM where it originates from the threat actor or an unconfirmed victim disclosure, and additionally marked "Threat-actor claim — independently unverified" where the source is FulcrumSec itself. ASSESSMENT is used for reasoned analytical judgement, always with an attached confidence level. UNKNOWN/"No reliable public evidence identified" is used deliberately wherever a requested topic could not be populated with genuine evidence, rather than left silently blank or filled with inference.

27.4 Attribution Methodology

The Manchester Airports Group attribution assessment in Section 4.6 was built by separately evaluating: (1) the victim's own official statements; (2) the threat actor's specific claims and the outlet's (BleepingComputer's) independent verification steps; (3) the consistency of the claimed technical method with the actor's previously documented, independently corroborated modus operandi; and (4) the presence or absence of any forensic, regulatory, or law-enforcement confirmation. Confidence was calibrated downward from "High" specifically because the victim has not confirmed the claim and no independent forensic finding exists, and upward from "Low/Unconfirmed" specifically because of the outlet's validated sample and the claim's technical consistency with the actor's known pattern.

About Shieldworkz

Shieldworkz is a specialist OT/ICS cybersecurity firm with an NDR solution and AI-based tools for securing SCADA, PLCs and Cyber Physical Systems. We serve critical infrastructure operators, industrial organisations, and government entities across the energy, oil and gas, manufacturing, utilities, transport, and defence sectors.

Our service areas include OT security assessments (powered by OThello Assess with sub-24-hour assessment cycles), NIS2 and IEC 62443 compliance programmes, OT threat intelligence advisories, OT SOC design and implementation, and regulatory readiness engagements for NCA (Saudi OTCC/ECC), NERC CIP, SOCI, and Singapore Cybersecurity Act obligations.

Reach out to us at info@shieldworkz.com for a briefing

NDR Solution: shieldworkz.com/products/ot-security-platform | **Media Scan:** shieldworkz.com/products/ot-security- | **Othello Assess:** shieldworkz.com/othello/assess | **Patch Management:** shieldworkz.com/products/patch-management-solution | **Regulatory Playbooks:** shieldworkz.com/regulatory-playbooks | **Remediation Guides:** shieldworkz.com/remediation-guides | **Reports:** shieldworkz.com/reports

About Shieldworkz

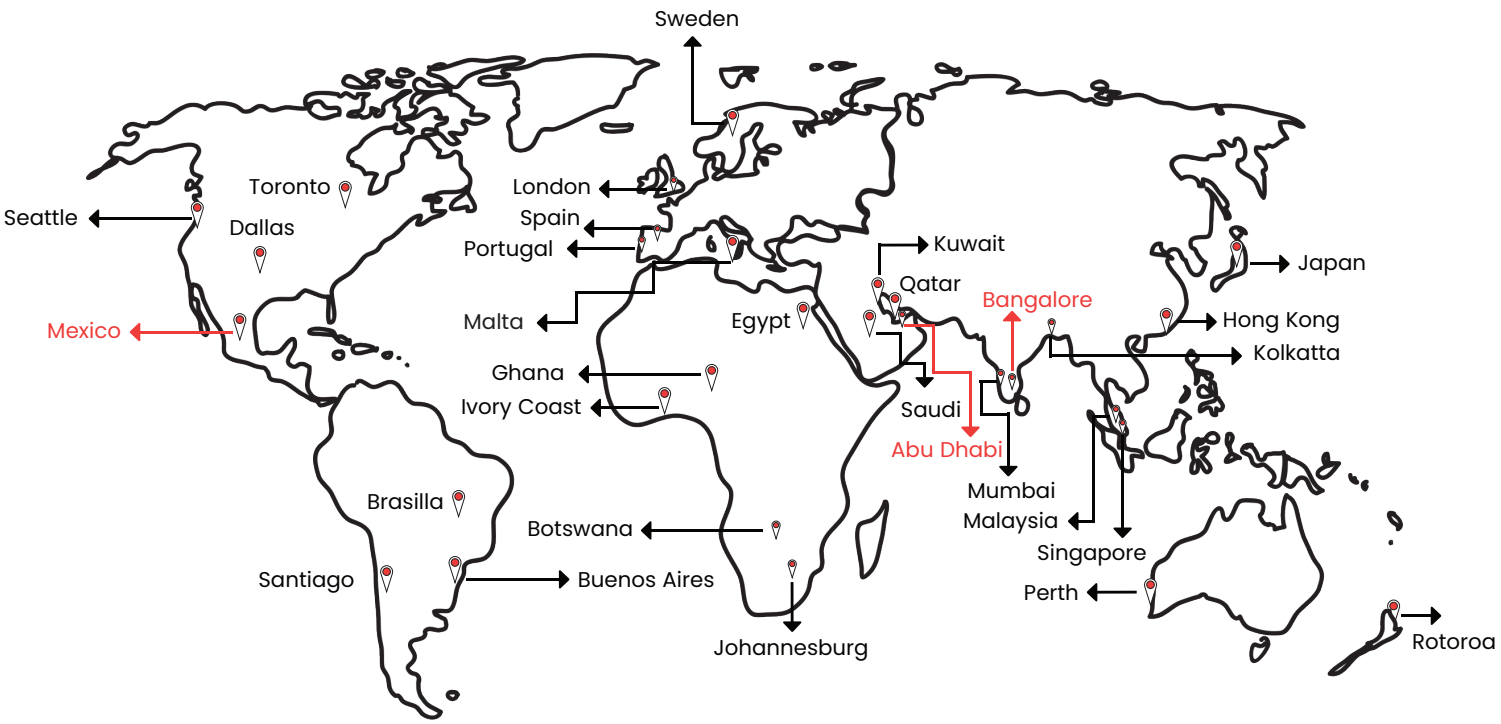

ISOC and Honeypot Locations

Honeypot Locations

Security Operations Center

Shieldworkz is a global OT security company founded by top industry experts to protect critical infrastructure using proprietary technology and a leading consulting platform, we partner with businesses to secure assets, networks, and programs across industries. Our services are tailored to each client's cyber risks and backed by the world's largest OT and IoT threat intelligence facility and a global research team.

Secure Your Industrial Future

Talk to us today!

From OT security assessments covering NIS2, IEC 62443, NERC CIP and other regional requirements to an OT security platform, Shieldworkz covers all compliance and industrial cybersecurity enhancement needs. Talk to us to learn how you can enhance your security posture in 7 easy steps.

Contact Us

6160 Warren Pkwy Ste 100
Frisco, TX 75034

+1 (323) 622-4481

americas@shieldworkz.com

396 Old Colony Road
Richmond Hill,
ON L4E 5A5 Canada

+1 (437) 223-7770

americas@shieldworkz.com

Tenth floor,

Abu Dhabi,
United Arab Emirates,
FAB Business Center

+971 56 660 5200

middleeast@shieldworkz.com

Fritz-Schäffer-Street 1,
4th floor

Bonn, 53113, Germany

+49 (0) 228 / 929 39210

europe@shieldworkz.com

Gopalan Signature Tower,
No 6, 2nd Floor, Old Madras
Road, Benniganahalli

Bengaluru,
Karnataka 560093

+91 9059620557

apac@shieldworkz.com